OPTIMAL OPERATING RULES FOR MULTI-RESERVOIR SYSTEMS

BY RICHARD MICHAEL MALES

SEPTEMBER 1968

DOCTOR OF PHILOSOPHY DISSERTATION

SUPPORTED BY

DEPARTMENT OF THE ARMY

GRADUATE FELLOWSHIP PROGRAM

OPTIMAL OPERATING RULES FOR MULTI-RESERVOIR SYSTEMS

BY

RICHARD MICHAEL MALES

S.B., Massachusetts Institute of Technology

(1963)

S.M., Massachusetts Institute of Technology

(1964)

Submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

at the

Massachusetts Institute of Technology

(September 1968)

Signature of Author . . . Richard Michael Males . . . .

Department of Civil Engineering, August 19, 1968

Certified by . . . .

Thesis Supervisor

Accepted by . . . .

Chairman, Departmental Committee on Graduate Students

Second Printing April, 1970

ABSTRACT

OPTIMAL OPERATING RULES FOR MULTI-RESERVOIR SYSTEMS

by

RICHARD MICHAEL MALES

Submitted to the Department of Civil Engineering on August 26, 1968 in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

The problem of determining optimal operating rules for multiple-reservoir systems is not presently susceptible to pure analytical treatment. The present study includes the operating rule as a parameter of the system, to be optimized by determining the appropriate form for the rule and appropriate values of the operating parameters. A simple policy, called the standard operating policy, is used instead of more complicated policies. A proposed technique called Linear Programming-Simulation-Search (LPSS) is investigated. The LPSS technique uses a linear programming model to obtain initial values for operating parameters. Starting with these values, simple search techniques of sensitivity, single-factor, and marginal analysis are used to determine near-optimal values of the parameters by searching the benefit response surface. The benefit response surface is defined through the use of a simulation model of the basin, which yields benefit values as a function of the various operating parameters.

A conceptual framework for the problem is presented, involving the definition of a Q set for the mathematical model. The Q set is used to define characteristics of the modelling process and operating rule, and conceptually represents the quality of the reproduction of reality in the mathematical model. Search processes to determine optimal operation also search over the Q set by varying the heirarchical level of the linear programming model and the type of operating rule used. The conceptual framework defines and clarifies the planning and decision-making process associated with water resource systems analysis, and provides a theoretical context for a discussion of searches in policy space.

Three different computer simulation models are used to test the feasibility of LPSS. Two of the models have associated linear programming models, and the relationship between the simulation and linear programming models is studied.

Results indicate the feasibility of the LPSS technique. A discussion of the appropriate type of policy and model for use of LPSS is presented. Releases determined by steady-state linear programming models are found to be consistently lower than near-optimal releases obtained by LPSS for stochastic simulation models.

Some basic problems with the approach so far are pointed out and methods of overcoming them in the future are recommended. The necessity for a data structure and data handling language for simulation and search processes is pointed out, and a form for such a language is suggested.

Thesis Supervisor: Ronald T. McLaughlin

Title: Associate Professor of Civil Engineering

ACKNOWLEDGEMENTS

The author is greatly indebted to Prof. R. T. McLaughlin of the M. I. T. Department of Civil Engineering, supervisor of this thesis, who originally stimulated the author's interest in water resource system analysis, and has carefully guided him through three theses.

Mr. David Hellstrom of the Corps of Engineers, New England Division, was more than liberal of his time and effort, and provided needed perspectives of real engineering situations.

Mr. Juan Antonio Poblete joined in fruitful discussion to examine various points presented in this thesis, and developed a number of the linear programming models used herein.

The research is part of a project supported by the M. I. T. Inter-American Program in Civil Engineering. During the course of this part of the research, the author was supported as a Graduate Fellow of the New England Division, Corps of Engineers, and is deeply grateful for this support and the many kindnesses extended by the personnel of the New England Division.

The computational work was performed in part at the Information Processing Services Center, M. I. T., Cambridge, Mass.

This thesis is dedicated to Redsie, Miski, Sarah, Meredith, Jane, and Mr. Fred Rogers, who all tried to provide some balance.

TABLE OF CONTENTS

CHAPTER VI - CONCLUSIONS AND RECOMMENDATIONS

CHAPTER I

INTRODUCTION AND PROBLEM STATEMENT

1.1 Introduction and Project Background

The investigations presented in this thesis were a portion of a research project on Systems Methodology for Water Resource Development, conducted jointly at the M.I.T. Department of Civil Engineering, and the Planning Center, University of Chile. This project, initiated in 1963, has included studies of the application of mathematical programming and simulation models of river basins for planning of water resource development, particularly in the context of data structures typical of developing countries.

The project was based on a real river basin selected for demonstration. Early studies investigated the feasibility of simple linear programming models of the basin, and determined the associated data requirements. The optimal operation of a single reservoir was examined through a combination of dynamic programming and simulation models. Later, the effect of social and political structures on engineering planning was recognized and investigated. Detailed simulation and mathematical programming models were developed. Optimal operation of a multiple-reservoir system was recognized as an important and unsolved problem. The present study concentrates on solutions to this last problem.

1.2 Problem Statement and Approach

Since the present investigation is a portion of a larger project in systems methodology, this particular investigation is within the context of the methods developed by the project as a whole. A brief outline of the proposed methodology is given below.

In the initial planning stages, when only rudimentary data is available, simple mathematical programming models are used to indicate initial capacities and sizes of structures in the basin. These simple models also serve to point out areas in which additional and/or improved data is required.

As further data is acquired, more detailed mathematical programming models are developed. In the final stages, detailed simulation models in the form of computer codes are constructed. The simulation models are used in conjunction with appropriately generated "synthetic" hydrologies to study the response of the system. Choices among different plans are based on comparisons of the results of simulation studies based on the plans.

In order to eliminate particular operating decisions as additional parameters in the search for final component sizes, the use of simulation requires that the operation of the reservoir system be optimal with respect to project goals for each plan. At present, no method has been developed to insure optimal operating decisions for a real system. For models which are simplifications of reality, optimal operation can be approached, and depending on the degree of simplification, achieved. Operation derived from such models is not necessarily truly optimal for a real system, but if the model is selected appropriately, the operation should be near-optimal. The general problem dealt with in this thesis, therefore, may be stated as the development of a method for near-optimal operation of a simulation model which is expressed in the form of a computer code.

For single reservoirs, approaches based on dynamic programming have been developed which yield optimal releases at each decision period for a simulation model. For multiple reservoirs with stochastic inflow however, no direct optimization technique has been developed. The approach investigated in this study uses operating policies for multiple reservoir systems which do not have an internal optimization feature such as linear or dynamic programming. Such policies, termed non-optimizing policies, can be characterized by the specification of certain parameters, such as target releases. Optimization of such non-optimizing policies consists of a determination of the optimal value of the policy parameters. Optimal values of the policy parameters are those values for which the operating policy determines decision variables leading to maximum benefits. The approach used in this study is to determine such values by performing search studies in the parameter space, rather than by purely analytical methods. Mathematical programming models are used as an integral part of the search process.

The research investigates a number of different policies and studies the applicability of simple search techniques. Linear programming models of the basins studied are used to provide guidelines for the search studies and to indicate appropriate starting values of the parameters characterizing the operating policy.

Three different simulation models have been investigated. Two of the models have associated linear programming models, of varying degrees of sophistication. The first simulation model represents an artificial basin, considered as a preliminary study to develop and test the feasibility of the proposed search techniques. A linear programming model which closely parallels the artificial basin provides initial values of the parameters of the operating policy. The second model used is a detailed simulation of the Maule river basin in central

Chile, for which a number of linear programming models exist. The third model is of a portion of the Connecticut River Basin in New England, and is included to examine the feasibility of search techniques where no mathematical programming model exists in conjunction with the simulation model.

The character of research in systems methodology is to determine the appropriateness of proposed approaches to decision-making. Therefore, the specific characteristics of near-optimal operating policies for the basins studied is not considered to be the result of this investigation. Rather, the results consist of conclusions as to the effectiveness of the proposed search concept, and indications of the proper role of such studies in basin planning.

1.3 Scope and Limitations

The concept of "optimizing" operation through search techniques is useful both in the planning process for proposed facilities, and in the operation of existing facilities. For the planning process, the simulation models used will reflect the different plans to be investigated, but each simulation model will be fixed with respect to capacities and sizes for that particular plan. For an existing system, the simulation model will represent the physical capacities actually present in the basin, but the operating policy parameters will be variables. Thus, in both the planning and operation stages, the simulation model is considered fixed, and the search is carried out over the operating policy parameters.

A major feature of the proposed technique is that it is not limited to a particular basin or class of basins. Similarly, the form of the operating policy used is not restricted, although certain forms are found to be more suitable than others.

A relatively simple operating policy structure is used in the first two models, which has advantages when used in conjunction with linear programming. As noted above, such policies are non-optimizing. When the state of the art is such that sufficiently general optimizing policies are developed, they will undoubtedly be described by some parameters, and the concept of searching over these parameters to determine the best "optimizing" policy is still valid. Improvements in the form of the operating policy can only lead to operating decisions closer to the hypothetical true optimal solution.

The study intentionally confines itself to the simpler operating polices and programming models, in order to determine how much useful information can be derived from them. In problems of this type, the degree of sophistication of the operating policy, programming model, and search techniques used, will depend upon the marginal increase in benefits as compared with the marginal increase in effort and cost required to obtain the improvements. The trade-off will be determined in individual instances. Similarly, the decision on when to limit the search process and accept the currently near-optimal policy as the best that can be obtained will be determined in each case.

1.4 Structure of the Thesis

Chapter I presents the context of the present study in a total research effort, defines the problem, and indicates the scope and limitations of the present work. Chapter II presents a conceptual framework for the research carried out, and surveys previous work in light of this framework. The preliminary model is described in Chapter III, with associated results and conclusions. Chapter IV deals with the Maule model, and Chapter V describes the Connecticut River Basin model. A summary and conclusions for the total study are presented in Chapter VI, together with suggestions for further study.

CHAPTER II

THEORETICAL CONCEPTS AND PREVIOUS WORK

2.1 Theoretical Concepts

2.1.1 Operating Policies

For the purposes of this study, an operating rule (operating rule and operating policy are used interchangeably) is considered to be a set of instructions which specifies in some manner all decision variables pertaining to flows and releases during any time interval when such decisions are made. The specification implies no ambiguity; the instructions provided by the operating rule are complete. Other instructions, such as the specification of maximum or minimum pool levels at given times, or the proportional division of total release, where that release is unspecified, are considered to be constraints, since they do not determine flows and releases directly, rather the flows and releases must be consistent with the restrictions imposed.

A distinction must be made between an algorithm which determines a policy, and the policy itself. A further distinction is made between a policy and the release determined by the policy. Three distinct hierarchical levels are present. The algorithm may use dynamic programming, linear programming, experience, or any other process, to derive the set of specific instructions which constitute the policy. The releases prescribed by the set of instructions will vary according to the values of certain parameters, which are termed parameters of the operating policy. Target release in any period is such a parameter. Specific values of the parameters of the operating policy, such as target releases, may be determined in a number of ways. Thus, a linear programming model may be used to yield appropriate values of target release.

The third heirarchical level is that of the releases. Releases are obtained from the policy instructions when those instructions are acted upon in accordance with the current state of the system. System state is not considered to be a parameter of the operating policy, rather it is the independent variable from which releases are determined according to the functional relationship expressed by the policy instructions.

The relationship among the factors may be indicated symbolically as follows, with F indicating a particular algorithm, f a policy, G_i the parameters of the policy, K the releases, and M a vector characterizing the system state:

$$\begin{aligned} F &\rightarrow f && \text{(policy derived from algorithm)} \\ f(G_i, M) &\rightarrow K && \text{(release derived from policy)} \end{aligned} \tag{2.1}$$

2.1.2 Definitions

After an examination of the literature on operating rules for multiple-reservoir systems, it was felt that no general terminology existed sufficient to describe and discuss the structure of water resource systems with associated models. The following discussion provides a conceptual framework within which the nature of operating rules can be defined explicitly.

Consider the river basin to be the basic unit about which decisions are to be made. Only certain sites within the basin are feasible for certain forms of development, such as the location of a reservoir, power plant, or irrigated farming area. Once the sites for all such feasible developments have been determined as well as the connections between them, with no reference to the scale of development at each site, the result will be called a configuration, denoted by C^*. Thus, a configuration indicates the various elements (reservoir, irrigation district, power-plant, treatment plant, pumping station, etc.) and the manner in which they are connected to form a directed network, by rivers, channels, penstocks, etc. For a given basin, there may be various configurations of interest, but since no magnitudes are specified, there will be one configuration which can include all other configurations if some of its elements or connectivities are at zero level.

A structure, P, for the basin is obtained from a configuration by the specification of all physical magnitudes for the elements, and a specification of the hydrology in the region. The hydrology is given by the hydrology vector, X, which may be considered to consist of one of a number of alternative specifications of the hydrology, such as the historical record, replicate sets of synthetically generated hydrologies, or parameters of the appropriate probability distributions, together with cross and serial correlation coefficients.

Magnitudes or levels of development of the elements are given by C_p, the physical magnitude set. Magnitudes of the physical magnitude set will in general appear in constraint relationships, of the form flow less than or equal to a magnitude at which flood damage occurs, or storage less than or equal to capacity. Thus, while physical magnitudes may be decision variables in come cases, they act primarily as a part of inequality constraint relationships.

Certain physical magnitudes pertaining to the structure do not fall into this category; for example, the turbine characteristics curve. These magnitudes are given by the technological conversion transformation, T.

The transformation T consists of the various transformation functions for conversion between physical quantities of the structure, as the transformation function for power from water at a given power station, or the production function for crop yield from irrigation water supplied. For purposes of simplicity, it is assumed that only the efficient points, on the production function, are included in these functions.

Thus, a structure P is obtained from the more general configuration C^* by specification of the hydrology vector X, the physical magnitude set C_p, and the technological conversion transformation T. A structure represents the existing physical reality in a river basin, and defines the basin capability for developing useful outputs from the hydrologic inputs.

The inclusion of social, legal, and economic constraints and benefit functions for the physical outputs yields a <u>system</u> S for a given structure. A system defines the total "reality" in the basin, including pertinent social, political, and legal requirements in force in the basin. These requirements may not be explicitly stated in the basin itself, but may be significant factors in engineering planning; for example, a political power structure may be favorable to one form of development. Social, legal, and economic constraints are included in the <u>SLE (social, legal, and economic) constraint set</u>, C_{sle}.

System outputs are considered to be physical quantities and probabilities, such as the increase in crop yield from irrigation, the power produced, or the reduction in probability of a given level of flood damage. These values make up the <u>output vector</u>, Y. A specific value of Y is denoted by $\bar{Y}$. The benefit function for gross benefit as a function of outputs is denoted by B(Y), and a specific value of benefits due to specific physical outputs is therefore obtained as $B(\bar{Y})$. Then a system S consists of a structure with a specified C_{sle} and $B(Y)$. It is the form of the benefit function that must be specified, for specific values depend upon specific values of $\bar{Y}$, as yet undetermined.

One further set may be introduced at this point, describing the "qualities" of the particular mathematical and simulation models used to represent the system. For instance, certain models of water resource systems take into account serial correlation, while others do not - an element of the qualities set Q serves to define this. Similarly for the other factors which may or may not be taken into account, such as stochasticity, number of time periods, and certain qualities of the operating rule, such as prediction of expected inflow. The Q set in reality serves to define the faithfulness of the model in terms of reality. While no absolute measure of the sophistication of the model is possible, the Q set can serve to rank different types of model in comparison with a hypothetical perfect model. The Q set is a useful manner of clarifying the nature of the operating rule and type of mathematical model as parameters of the system.

After the system S and associated elements of the Q set have been designated, it becomes possible to construct a simulation model of the system, complete except for the operating rule, that is, lacking only instructions for arriving at decision variables in each period, in order to function. If the simulation is in the form of computer codes, then the simulation may be constructed with a call to a subroutine defining the operating policy at each point where a release must be specified. The state variables are supplied as data to the operating policy subroutine, which returns values for the releases. A simulation model constructed in the foregoing manner will be called a simulation model with an external operating rule.

Certain additional definitions will be made to aid in the description of the functional relationships between the system and the operating rule. For simplicity, let both C_p and C_{sle} be combined to form the single set C. Let $\bar{X}$ be a specific value of the hydrology vector X, $\bar{C}$ a specific value of C, and $\bar{T}$ a value of the technological transformation T. Define a release vector, K, to consist of all the releases necessary in any period to specify the operation of the system, and let $\bar{K}$ be a specific value of K. Let H represent the conversion function which maps the hydrologic inputs and constraints into the specific value of the outputs, $\bar{Y}$. Then the functional relationship between input, output, decisions, and constraints, is:

$$\bar{Y} = H(\bar{K},\bar{X},\bar{C},\bar{T}) \tag{2.2}$$

If the hydrology is specified in a unique fashion, the output is a single-valued function of the input, constraints, and decisions.

The value of the release vector is obtained from the operating policy. Let L_i be the ith such policy of the many possible policies. Then the releases are determined by the operating policy as:

$$\bar{K} = L_i(\bar{X},\bar{C},\bar{T},B(\bar{Y}),\bar{Q}) \tag{2.3}$$

where the parameters of the operating policy, G_i, are included in the Q set. In the global sense, i is a decision variable determining the form of the operating rule, and Q and L_i are linked and must be consistent. For a given Q, only certain policies are feasible, and for a given L_i, only certain Q are appropriate.

It is obvious that the formulation is implicit, since the two functional relationships state that the output is based on the input and the decisions, and that the decisions are based on the input and the form of the benefit function of the output. Consideration of a dynamic programming algorithm will show that a formulation of this sort is feasible. For this algorithm the decisions that will eventually determine outputs are made by a consideration of the form of the benefit function for these outputs and the consequences of using any set of decisions to obtain these outputs.

In practice, the conversion or transfer function H will be determinate from system properties, and in fact the construction of a simulation is equivalent to, and often the only manner of, defining H. Then, a given L_i may be chosen, yielding $\bar{K}$ from equation (2.3), the $\bar{Y}$ may be obtained from H by equation (2.2), and finally, $B(\bar{Y})$.

2.1.3 Sampling

As yet, no method has been found to lead directly to an optimal system, optimally operated, from a knowledge of X, C, T, and B(Y). As a consequence, many systems are usually investigated, and used to map a response surface. Search processes over this response surface attempt to determine the global optimum. One of the purposes of this study is to extend these techniques to the enlarged model which treats the operating rule as a parameter, and contains a Q set characterizing the sample point additionally.

The sample space E consists of those points defined by X, Y, T, and K, and satisfying the constraints implicit in C_p and C_{sle}. All such points are feasible along the surfaces defined by the input-output relation H. The inclusion of the L_i as decision variables increases the sample space and divides it into sub-spaces determined as particular

$\bar{X}$, $\bar{Y}$ combinations obtained with any particular L_i. Clearly, not all feasible $\bar{X}$, $\bar{Y}$ combinations can be obtained from all operating policies.

A true global optimization policy, L_{opt}, may be postulated, that yields $\bar{Y}$ such that $B(\bar{Y})$ is maximized, while satisfying all constraints and taking into account the most sophisticated model possible in terms of the Q set. The L_i may be termed boundary policies, because some features of the Q set are taken into account and may be thought of as having zero value. The use of many L_i for a given system is similar to the classical search techniques along the boundaries of the sample space. By sampling on many L_i, the L_{opt} rule is approached, or at least delineated in some of its characteristics.

The L_i partition the sample space E into sub-spaces, each associated with a Q set. Within a sub-space, the parameters of the operating rule may be varied to locate the optimal rule for that sub-space. This rule will achieve a local optimum on the response surface, but may not reach the global optimum given by L_{opt}. As a consequence, the achievement of near-optimal operation by search among the various local optima is the desired goal. The search is carried out over the form of the operating rule to determine those rules which yield the highest benefit regions in their associated sub-spaces. When promising rules have been selected, the parameters of the operating rule may be varied to define the local optimum for the rule, and the final selection made on this basis. True optimality cannot be achieved because the Q sets do not reproduce reality, and beacuse the search effort will not be able to investigate all Q sets and all parameter sets of the objective function. Due to the inherent nature of mathematical modelling, near-optimal operation is usually the only achievable goal.

The engineering decision-maker is usually forced to choose between points in the sample space according to their associated values of C and $B(\bar{Y})$, since in most cases all objectives cannot be expressed in economic terms in a suitable manner. The rule L_{opt} would then yield the maximum $B(\bar{Y})$ for any C. Since this rule is unknown, rules which give maximum $B(\bar{Y})$ for C and Q are sought. The sampling process may be considered to take the following form:

1) Numerous mathematical models, of varying detail, completeness, and accuracy are available, which yield various forms and parameters of operating rules. For a given B(Y), X, C, T, these models yield the release vector $\bar{K}$ and have associated with them the set Q.
2) A simulation of the system is considered to be available, lacking only the vector $\bar{K}$ for each period. The insertion of the various rules obtained from the models into the simulation yields values of $B(\bar{Y})$. The sampling process discovers those L_i and associated Q which lead to near-optimal points of E, as determined by the results of the simulation and a suitable search procedure.

The above is a formalization of the considerations to be used in discussing water resource systems with associated models. Clearly, the refinements are not indicated, such as a complete listing of the elements of the Q set, and a detailed formulation of this set. The above ideas serve essentially to conceptualize and organize the main features of the problem, and provide for a common language, sufficiently explicit to discuss the problem at hand without ambiguity.

2.1.3 Simple Search Techniques

When a form of operating rule has been selected, the problem is reduced to the specification of the best values for the parameters of the operating rule. This involves the delineation of the response surface, and a search for the optimal peak. A number of possibilities exist for locating the optimal peak. Most techniques presented describe a method of improving the position on the response surface, starting from an initial position. The initial position can be determined in a number of manners.

Uniform or random grid sampling provide a number of points from which to select an initial value for "hill-climbing". Linear programming can be considered as a search technique that determines a single set of operating policy parameters. This single set provides a single point on the response surface, from which improvements can be made. Thus, a search process consists of two stages. First, an initial point or points are selected. Secondly, the position on the response surface starting from these points is improved.

Single-factor analysis is one method of indicating an appropriate direction from a given point. Each of the parameters is varied by a fixed increment, both increasing and decreasing the initial value of the parameter. The parameters are varied one at a time, with the remaining parameters maintaining their initial values. Therefore, the number of tests performed is equal to twice the number of parameters. This procedure serves to indicate the gradient of the response surface in the direction of each of the parameters. In many cases, the response surface will be flat in some directions, indicating indifference to change in that direction, and thus temporarily eliminating the particular parameter from consideration. The advantage of single-factor analysis is that it yields maximum information about the nature of the response surface in the region of a point.

Axial analysis is used in this study to indicate a search along the axes of the policy parameter space, starting from the origin, and varying each parameter individually to obtain a profile along its axis. Any test in an axial analysis maintains all parameters except one equal to zero. The number of tests depends upon the grid spacing and range of interest for the values of the variable. The advantages of axial analysis are that it provides an indication of the shape of the response surface, and that it can be used to define the limits of further studies, by noting those regions along the axis where benefits decrease sharply as the parameter is increased along the axis. Certain variables will show a horizontal profile at sufficient distance from the origin, particularly those variables which are limited by availability of water. Thus, increasing the target draft for a reservoir far beyond the available water will not result in any change in benefits, since this indicates that the parameter has gone beyond the range of physical feasiblity, i.e. the indicated draft can never be supplied with the existing hydrology.

Axial analysis can also be used to obtain a starting value for hill-climbing by selecting the value of each parameter that maximizes benefit along its respective axis, and combining these values into a single parameter set. Axial analysis is simple and straight-forward and need be performed only once for any search process.

Marginal analysis is used here to indicate an analysis performed when all but two variables have been fixed. Then all combinations of values of the two variables within a range of interest are examined for fixed values of the other parameters. The number of tests depends upon the range of interest and grid spacing, and is the product of the number of test values selected for each variable. Marginal analysis is good at bringing out the fine structure of the response surface, since it places a grid over the two variables of interest. After the optimal combination of the two variables has been determined, another set of variables can be examined in this manner, thus optimizing "two at a time". This is a substantial reduction in effort from placing a grid over the entire response surface. Marginal analysis is particularly useful if the parameters naturally fall into pairs, as in the case for a two-period model.

Finally, sensitivity analysis is used here to indicate a study determining response to changes in a single variable, with all others being held fixed. A single variable can often be isolated and fixed by sensitivity analysis at an optimal value during an early stage of the search process, reducing the sample space.

Engineering "feel" and familiarity with the model are additional tools that are important in eliminating parameters and indicating promising regions. Occasionally, a person who has constructed a simulation model will have acquired sufficient familiarity with the model to predict possible performance. Results of a test may be analyzed with a thorough concept of what has taken place physically in the simulation, rather than treating the simulation as a "black box" which yields a benefit value as output for a certain set of inputs.

Search techniques have been restricted to the simple methods noted above, both to reduce the scope of the investigation and to investigate the applicability of the simpler techniques. Optimizing multi-variate search techniques may more rapidly identify optimal regions, but as always, trade-offs between such benefits and the efforts involved in establishing a suitable optimizing multi-variate search procedure must be considered.

2.1.4 State and Release Oriented Policies

One simple division of operating policies into two separate classes, which would be reflected by two different values of a parameter in the Q set, is the distinction between state and release oriented operating policies. For the purposes of this study, a release oriented rule is a rule which yields the release from a reservoir directly as a function of the input parameters for the operating policy, without considering variations in the releases due to expected differences in inflow, high reservoir storage, etc. A release oriented policy would attempt to meet a target draft under all circumstances, and the target draft would not be revised or changed depending on the amount of water available.

A state oriented rule would modify the target drafts according to the water currently available. Thus, a flow higher than expected in a given month would allow for a release greater than the usual release during that month. Similarly, a deficient flow would modify the target release by reducing it somewhat, rather than exhausting the reservoir storage by attempting to supply the normal release during the month. A state-oriented policy needs additional parameters beyond target releases in every period, since some estimate of "normal" flows and typical volumes is needed to determine if water currently available is normal, excess, or deficient. A linear programming model is particularly suitable for yielding information about both target releases and normal flow and storage levels. Thus, an attempt to utilize linear programming and simulation models in conjunction would consider using the data obtained from a linear programming model to indicate target drafts, modified according to whether or not the reservoir state corresponds to the state predicted by linear programming.

With any sort of rule that modifies decision variables based on data, the decisions are only as good as the data. Thus, a state oriented rule, which is expected to be superior to a release oriented rule, since it modifies fixed releases in accordance with current conditions, may in fact prove inferior to the fixed rule, if the predicted values of state variables do not correspond with reality. Adaptive control can be built into a state oriented rule, to revise initally poor estimates, and such adaptive control is certain to be a feature of any true optimizing policy which must make predictions or compare with "normal" values for flow, storage, etc.

2.2 Survey of Previous Work

Although there is at present a fairly substantial body of work dealing with various techniques of water resource systems analysis, material treating directly the problem of optimal operation of multiple-reservoir systems is not so readily available, reflecting the fact that this is one of the principal unsolved problems in water resource systems analysis.

Maass, Hufschmidt, et.al., in Design of Water Resource Systems (1), presented the first large-scale study of systems analysis techniques for water resource systems. A description of a detailed simulation model is presented, including features of the computer programming. The operating rule is non-optimizing, with parameters specifying target releases and flood-storage capacities. The Q set does not include prediction of future inflows. The parameters of the operating rule are not clearly identified as such.

Hufschmidt (2) discusses sampling techniques for examination of the response surface generated by the simulation model, but does not treat the operating policy as a parameter. Dorfman (3) describes a simple 2-period linear programming model for a basin with two reservoirs in series, and a Q set specifying predictable hydrology. The complicated method of treating the non-linear objective function is no longer appropriate due to advances in mathematical programming techniques. The extension to more periods is demonstrated, and another model in which the probability distribution for inflow is given is discussed. These constitute increasing levels of sophistication of the Q set.

Thomas and Watermeyer further "upgrade" the level of the Q set by using stochastic linear programming (4) to determine the shape of the optimal policy for a two-period model of a single reservoir. They indicate that the technique is restricted to single reservoirs, and does not take into account serial correlation of the inflows.

Wallace (5) develops a number of linear programming models of the Maule basin, with varying Q set, including a 12-period monthly model with predictable hydrology. He also demonstrates a method of handling serial correlation with stochastic linear programming, eliminating a restriction in the work of Thomas and Watermeyer. Computational restrictions limit the practicability of the technique.

Parikh (6) proposes a method for optimal operation of multiple reservoir systems, assuming complete knowledge of inflows over the planning period. His method requires that individual reservoirs be optimized with respect to total system goals by dynamic programming, and then indicates an iterative technique for global optimization of the system. Operation of a 4-reservoir system over 10 years is optimized for the known flows.

Hall, Butcher, and Esogbue (7) utilize the dynamic programming technique for a single reservoir that is postulated in Parikh's work, again using complete knowledge of inflows. The operation of Shasta Dam for firm and dump power and water yield was optimized by this technique, with reasonable computational effort.

Bower, Hufschmidt, and Reedy (8) derive generalized operating procedures to accomplish certain physical objectives, such as the minimization of future spill. These procedures are not associated with optimal operation for system goals except in so far as they conserve water or distribute losses over time. A general discussion of the nature of flexible operating rules is presented.

Hufschmidt and Fiering (9) describe a detailed computer simulation model of the Lehigh River Basin in Pennsylvania. An operating policy making use of the space rule, described in (8), is developed. The policy is non-optimizing and characterized in part by parameters. The parameters are varied along with others to optimize total system design. The concept of optimization of operating policies in conjunction with simulation models is presented, but the operating parameters are not disassociated from other system parameters in performing the studies. As a consequence, there is no assurance that the appropriate "near-optimal" operating policy is being used for any values of the system design components, and the operating policy remains a parameter of the model. As in the simulation model in (1), the operating parameters are not clearly characterized as such.

A dynamic programming optimization technique is presented by Schweig and Cole (10) for operation of two reservoirs in parallel supplying the same demand, principally to demonstrate the feasibility of such a technique. Simplifications and restrictions limit the usefulness of the results.

Optimal operation of single reservoirs exclusively is treated in a number of papers. Little (11) optimizes monthly releases for hydropower in Grand Coulee Dam using dynamic programming in one of the earliest treatments of the subject. The Q set did not include serial correlation among inflows, but specifies a probability distribution for hydrology rather than a predictable hydrology. Buras (12) treats a reservoir and ground water basin used conjunctively, by dynamic programming. The resulting optimal policy is state-oriented, and is presented in tabular rather than algorithmic form.

Hales (13) uses dynamic programming to develop an optimal operating rule used in a simulation model of a single reservoir, and investigates the effects of forecasting of future inflows and reliability of streamflow data on net benefits. Some search studies are performed over parameters of the operating rule to choose optimal values, and the elimination of the operating rule as a parameter is recognized. Young (14) uses a similar approach, but develops a forward-looking dynamic programming algorithm. The optimal policy is compared with a fixed policy, and is found to be a significant improvement. A mathematical formulation of the single-reservoir optimization problem is also included.

Ibanez (15) has done a survey of various types of operating policies, and developed computer codes for four different policies. Little testing was done, but preliminary results indicated that a policy based on data from a linear programming model yielded highest benefits on a 10-year simulation run. Ibanez tested his operating policies on a detailed simulation model of the Maule River Basin in Chile, developed by Leonvendagar and McLaughlin (16). The work of these investigators is drawn upon substantially in Chapter IV of this report.

Poblete (20) has developed more detailed linear programming models of the Maule basin than those of Wallace. Characteristics of 2, 4, and 12-period models of the same basin are investigated. Results of his work are used in this study to define target parameters for operating policies in the Maule basin.

Finally, Hellstrom (18) has developed a generalized physical simulation model and is currently using it in studies of the Connectictut River Basin. Investigations using this model are presented in Chapter V.

In general, the literature survey indicated that the majority of researchers have not adequately located the optimal operation problem in the spectrum of total project planning. Techniques and methods proposed do not lead to useful planning decisions. There appears to be little recognition of the relation between reality and modelling, and for this reason many studies concentrate on optimization of the model performance as the goal, rather than optimization of system performance. A clarification of the precise goals involved in the studies is required, and appropriate interfacing with other techniques has been overlooked. There appears to be a trend towards development of more and more sophisticated models, without attempting to determine how well the simpler models perform, and how appropriate the sophisticated models are in the light of existing data structures.

CHAPTER III

THE PRELIMINARY MODEL

3.1 Introduction

This portion of the study constitutes an investigation of the feasibility and characteristics of the proposed search technique, which will be called linear programming-simulation-search (LPSS). An LPSS process is used to define "optimal" values of operating policy parameters in the following manner:

1) A linear programming (LP) model generates initial values of the operating parameters, consistent with the form of operating policy to be used in the simulation.
2) The initial values of operating parameters developed by the linear programming model ared used in a simulation to obtain a value of benefits corresponding to the linear programming parameters.
3) Parallel to the development of the LP model, random or uniform sampling and axial analysis are used to generate other parameter sets, and simulation runs determine the associated benefits.
4) An examination of the results of Steps 2 and 3, combined with insight into the characteristics of the simulation model, is used to determine a best initial value.
5) Single-factor, sensitivity, and marginal analysis are used as deemed appropriate during the search process to improve the values of the parameter set. Axial analysis is used to delineate the boundaries of the response surface.
6) The search process continues by the methods of step 5, until

the trade-off between further search effort and improvement in benefits dictates a halt, at which point the currently "near-optimal" parameter set is accepted as defining the operating policy.

The use of LPSS requires the existence of linear programming and simulation models of the basin of interest. The form of the operating policy used in the simulation model should be consistent with the type of data obtainable from an LP model.

Such models were constructed, with attention to the minimum coding effort that could go into a simple simulation model designed to test LPSS. The simulation model was coded in Fortran IV and implemented on the IBM 360/40 and 360/65 computers, while the linear programming model was generated and run on an IBM 1130 computer, using the LP-MOSS 1130 Linear Programming System (17).

3.2 The Simulation Model

3.2.1 Structure of the Simulation

The configuration for the preliminary model consists of 4 reservoirs, 2 power plants, two independent inflows, and a downstream irrigation zone, connected as shown in Figure 3.1. This configuration contains as a subset the simpler cases of parallel and series reservoirs, and is sufficiently sophisticated in itself to present some interest. In addition, if techniques for decomposition of reservoir systems into subsystems are developed, then the configuration chosen represents the minimal subsystem which includes combinations of parallel and series reservoirs, such that all larger systems can be built up from parallel and series combinations of this subsystem. This applies only to the

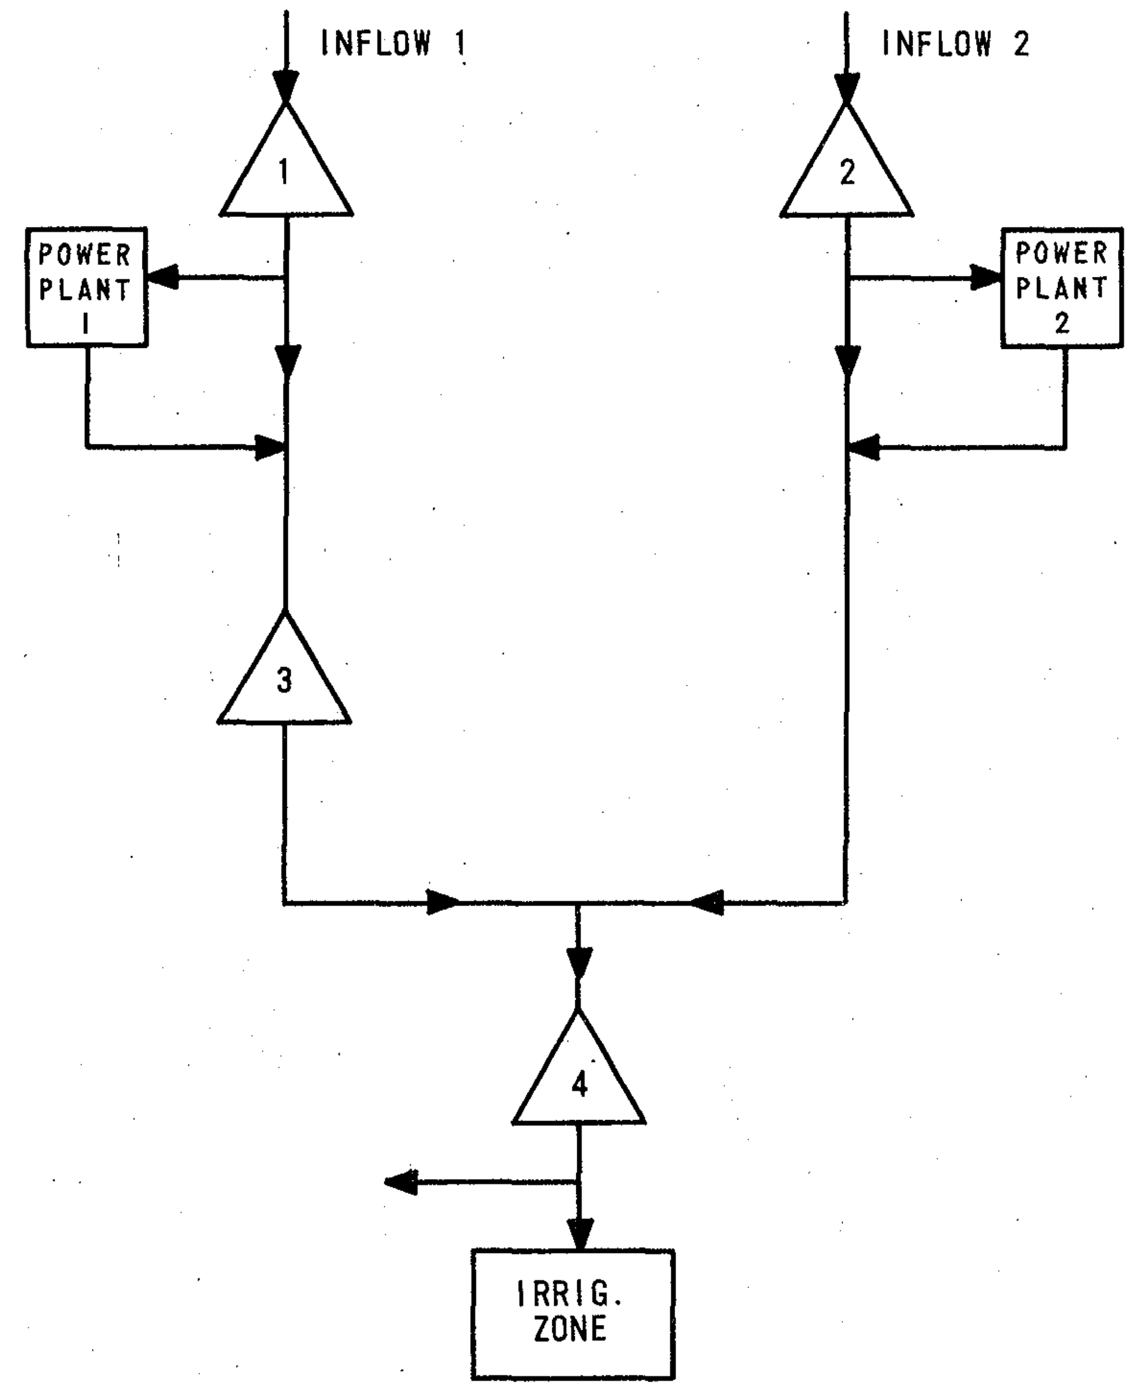

FIGURE 3.1

PRELIMINARY MODEL

reservoir configuration, not to the topology of the other elements.

A generalized simulation model for this configuration was developed and coded, with external operating rule and benefit functions. The model is discrete, with the number of periods per year a variable set by the user, between 1 (a yearly model) and 12(a monthly model). The periods should be of uniform length. As coded, the model provides for a maximum simulation run of 500 periods, or about 42 years for a monthly model.

The hydrology of the two rivers in the configuration is supplied to the simulation on punched cards, one card per period with two values per card, for the total number of periods of simulation. The units are arbitrary volume units per period, such as m^3/month, or acre-feet/season. Reservoir capacity for the four reservoirs is given in the same volume units as for inflow, and power plant and irrigation zone capacity are similarly parameters of the simulation given in volume units per period. Reservoir and power plant capacity are constant for all periods, but irrigation capacity may vary with the period, to reflect differences in crop water requirements from period to period.

Additional parameters of the simulation are the interest rate per period, and the initial volume of the reservoirs at the start of the simulation period, in consistent volume units.

The operation of the simulation model proceeds as follows. The hydrology is read into a matrix, from which values are taken one at a time during each period for each river. The parameters of the benefit functions, and capacities of the elements, as well as the interest rate are read from punched cards. At this point, the operating parameters corresponding to the particular simulation run are read. The simulation initializes all parameters to their appropriate values, and commences the operation of the first period by calling the operating rule subroutine. This subroutine will be described in detail below. The operating policy subroutine determines appropriate values for releases from the four reservoirs and diversions to the two power plants and the irrigation zone. Control is returned to the main program, which revises the state of the reservoirs according to the releases as determined by the operating rule, and proceeds to call the power benefit subroutine, which is called each period.

Power benefits are calculated in arbitrary benefit units as a function of the diversion to each power plant, in each period, and the shape of the functional relationship may vary from period to period over the year. The form of the function, however, is constrained to be two-piece piece-wise linear by the coding of the power benefit subroutine, but the slopes and intercepts are user-specified parameters in each period over the year. Corresponding periods, or seasons, in different years, will have the same benefit function.

For the present simplified model, the power benefit is a function only of the volume of water supplied to the power plant. No account is taken of the technological transformations between flow, head, efficiency, and power at the plant. For greater sophistication, a more detailed evaluation of power production and benefits could be handled by a more refined subroutine, but for the purposes of the preliminary model, the discharge-based form of the benefit function proves entirely adequate.

The power subroutine calculates benefits from the two-piece function, and discounts benefits to present value at the start of the simulation period, then returns to the main program, where a running total of the power benefits is kept.

If the current period is the final period of the year, then the irrigation benefit subroutine is called to calculate irrigation benefits as a function of total irrigation diversion during the year. As in the power benefit case, the irrigation benefit is in arbitrary benefit units, derived from a single two-piece piece-wise linear function, with the exception that it is a single function for the entire year, and does not vary from year to year. Irrigation benefits are taken as functions of the total yearly diversion rather than the diversion per period to take into account in a simplified manner the nature of irrigated growing. No benefit is provided until the crops are harvested, in contrast to power benefits, which may be considered to be of value for sale continuously throughout the year, although the price may vary.

The operating policy contains a provision for a maximum irrigation diversion in each period, thus insuring that a large total for the year cannot be obtained by supplying an excess in one period and a deficiency in a later period or vice versa, both situations being detrimental to the crop yield.

It should be noted that the benefit functions are considered to be net benefit functions. There is no calculation of capital cost and operating costs are already subtracted from gross benefits. The situation simulates the case of constructed and paid-for facilities with minimal operating cost compared to benefits.

As in the case of power benefits, the irrigation benefits are discounted to present value, and control is returned to the main program, where a running total of irrigation benefits is kept. The procedure starts again with the following period with another call to the operating policy subroutine, and so on until the entire simulation run is completed. At this point pertinent information is printed out, and the parameters for the next run are read in and the process repeats for the desired number of runs.

Output from the program is available each period if the user desires. In this case, at the end of each period the program will print out the complete state of the system, including reservoir volumes, releases throughout the system, inflows, and power and irrigation benefits for the current period. At the end of the run, the summaries and totals are printed out as before.

3.2.2 The Operating Policy

As noted previously, the operating policy subroutine for this simulation is "external". That is, except for input/output functions, which are included in the main program, the operating policy may take any form whatsoever provided that at the end of each subroutine call, unambiguous releases at the pertinent points throughout the basin are generated. As with the benefit functions however, a particular form of operating poicy was most suited to the purposes of this study.

This form makes use of the so-called Standard Operating Policy, discussed by Fiering (19) and Young (14) for single reservoirs and represented graphically in Figure 3.2. An examination of this figure shows that the policy may be represented by a piece-wise linear graph determined by two parameters, the reservoir capacity and the target draft in the period in question.

Releases are determined unambiguously for a reservoir by this policy as follows. Initially, the water available for use in any period is determined. This is often taken to be the stored volume at the beginning of the period plus the actual or expected inflow during the period; the

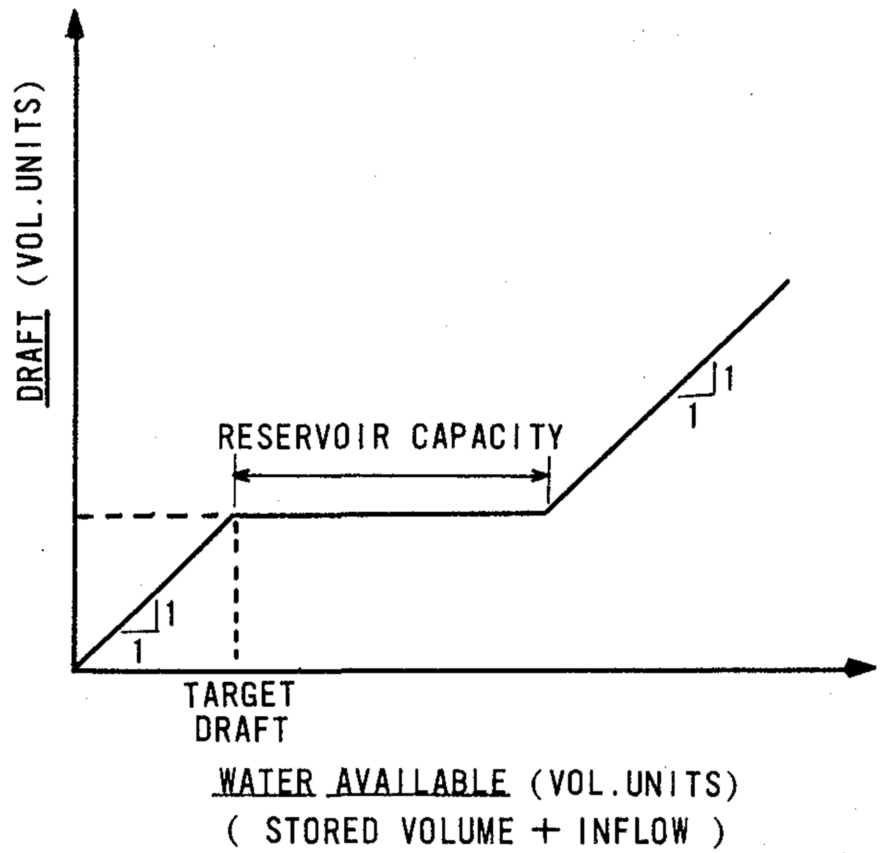

FIGURE 3.2

STANDARD OPERATING POLICY

model under discussion uses stored volume plus actual inflow. This quantity of water is compared with the target draft for the period. If it is less than or equal to the target draft, then the entire amount of water available will be supplied. If the water available is greater than the target draft but less than the sum of target draft plus reservoir capacity, then the target draft will be supplied, and the remainder of water available will be stored in the reservoir. As the volume of water available increases still further, a point will be reached at which the target draft can be supplied and the reservoir filled to capacity. Any additional water must then be spilled. In essence, the policy operates a reservoir to meet committments if possible, then to store as much as possible up to capacity.

This form of operating policy can be particularly useful in simple simulation models due to its simplicity, the ease with which it can be coded, and the fact that it depends on only one parameter other than the reservoir capacity, which is fixed for a given study. Further, a value for target draft in a period is very readily obtained from a simple linear programming model such as the one presented in section 3.3. Thus, the operating policy parameters are limited to one parameter per reservoir per period.

Early results of Ibanez (15) indicated that an operating policy for multiple reservoir systems based on the standard operating policy with target drafts arrived at by a linear programming model developed by Wallace (5) yielded higher net benefits over a 10-year simulation run than three other types of policies tested. As a consequence, this form of operating policy was selected for study in the preliminary model.

Releases are determined sequentially starting at the reservoirs furthest upstream, for which the water available is taken as the sum of inflow during the period plus stored volume at the start of the period. The standard operating policy is then used to determine releases from the upstream reservoirs. Operation proceeds downstream, with the water available being taken as the sum of stored volume plus water released from the upstream reservoirs. This is clearly a simplification of reality, in that such factors as transit time, local inflow, channel losses, etc., are not taken into account. Including these factors only implies an increase in sophistication of the model, and does not require the introduction of any new techniques.

The operating policy thus proceeds sequentially downstream, determining releases. Diversions to power plants and the irrigation area are treated in a similar manner. A target value is determined as input data to the operating rule, and this target value may be determined by a linear programming model. In any case, the target values must not be greater than the capacity of the facilities during the period in question. The quantity of water available at the diversion point is compared with the target value, and water will be diverted until the target value is reached. At this point, no further diversion will take place. Thus, the operating policy for the facilities has the form shown in Figure 3.3.

A typical operating procedure used for hydroelectric plants involves generation of firm energy at a normal target value below plant capacity, and generation of additional dump or off-peak energy up to plant capacity when water is available. Typical benefit functions provide a certain value for meeting the target power, a loss function for not meeting the target value, and an incremental value for dump power less than the unitary value of energy produced at the target level. The present model

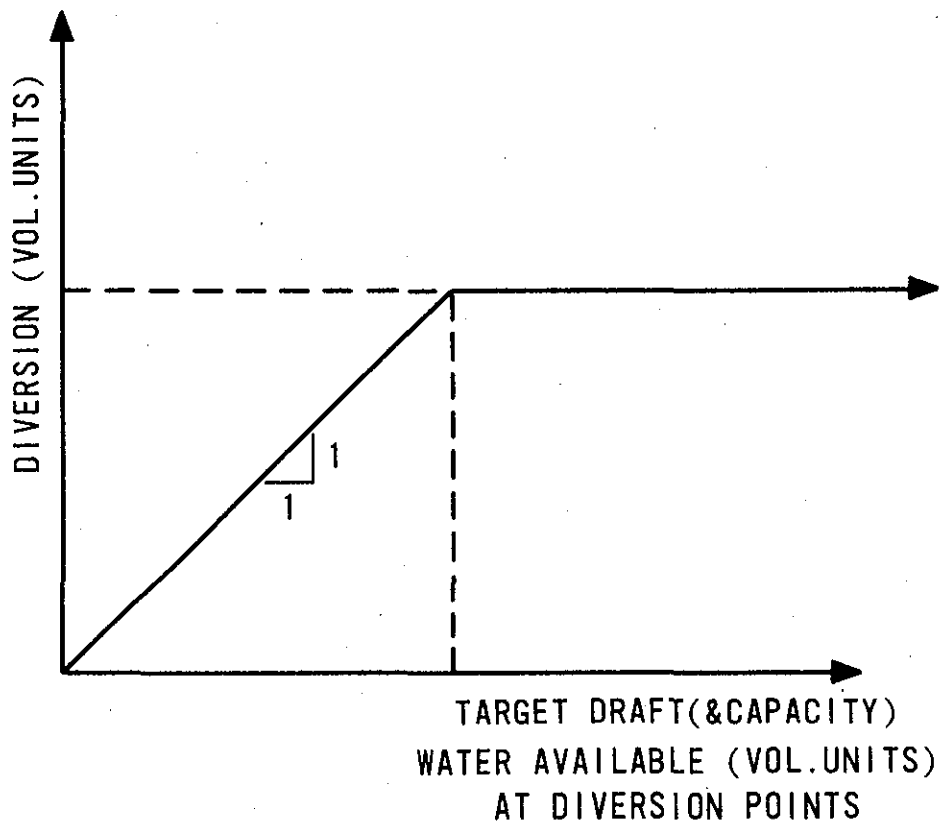

FIGURE 3.3

STANDARD OPERATING POLICY FOR FACILTIES

does not take into account production of dump energy above a target level, and the power benefit function is linear to plant capacity, with no loss function associated with not meeting the target value. Accordingly, such an economic structure dictates that the appropriate target diversion to power should equal the plant capacity. Falling short of this target does not incur an increasing loss, as would be the case with a typical non-linear power benefit function, and setting the target diversion to capacity insures that all available water will be used for power generation up to plant capacity.

An examination of the structure of Figure 3.1 shows that, with the given linear benefit function, this is an efficient operating procedure for the power plants. If a power benefit function including losses is used, or if return flow from power plants may be diverted to another basin or sub-basin, or if changes in water quality must be taken into account subject to total basin constraints, setting target diversion to plant capacity will not necessarily be efficient. For the present case, however, the targets may be set at capacity.

3.2.3 Data for the Simulation Model

For operation of the simulation model, hydrology, interest rate, size of reservoirs, power plants, and irrigation district, irrigation and power benefit functions, and initial volumes must all be specified. In the terminology developed in Chapter II, the simulation model coding specifies the configuration, and the particular structure is determined by specification of X and C_p. In this simplified model, no T transformation is required, since all parameters are in common flow and benefit units, and the transformation is assumed to have been made externally in the construction of the benefit functions.

At the start of this stage of the research an attempt was made to locate an actual basin with the desired configuration of elements, and with appropriate data available. It soon became obvious that selection of an actual basin would require efforts in data collection and organization which were not pertinent to this part of the research. As a consequence, an artifical data structure for the configuration was developed, with the actual data values selected by the author. This served to concentrate effort on the study of the LPSS process itself, and is not considered a fundamental limitation since Chapters IV and V deal with basins with real data structures.

For the model, the X vector was chosen as an explicit specification of hydrology by flows in each period over the simulation duration. The word "hydrology" is used here to indicate this kind of specification, rather than other alternative possibilities.

Initial testing and development of the model was performed using a 5-year two-season hydrology. A 50-year hydrology was generated from this 5-year hydrology by duplicating the punched cards and shuffling them, maintaining proper seasonal order. The mean values of the 5-year set were used as input to the linear programming model, and the initial 50-year hydrology, termed data set 2, had identical mean values by the method of generation. Further, the generating method reduced the influence of stochasticity in creating differences between the simulation and linear programming models. Search techniques were tested and developed using data set 2. For further studies, six synthetic hydrologies were developed from a Markovian recursion model using a normally distributed variate, with mean and standard deviation equal to that shown in Table 3.1. Regression and correlation were not taken into account in the generation of the six "synthetic hydrologies".

Capacities of reservoirs, power plants, and irrigation district were selected in accordance with the hydrology, so that the power and irrigation demands would present a competition in use, with the demands slightly greater than the mean inflow for the year, and out of phase with the seasonal inflows. Prior to the development of the linear programming model, and for the initial studies, reservoir capacities were selected large enough to hold at least the mean upstream inflow for the year. Table 3.1 lists the pertinent data values.

Benefit functions, $B(\bar{Y})$, for the model, are two-piece piece-wise linear functions to element capacity. The general form is as shown in Figure 3.4a. For the simulation model to coincide with the linear programming model, however, the benefit functions for power and irrigation are linear to element capacity, with target output equal to element capacity, as shown in Figures 3.4b, c, and d.

Interest rate and initial volume of the reservoirs at the start of simulation were taken as parameters whose influence was to be determined in the initial stages in order to fix on appropriate values for the LPSS procedure. Further studies were undertaken to examine the sensitivity of the optimal policy to these factors and to reservoir capacity. These studies will be discussed in a later section.

Further data for the simulation model consists of target releases at the four reservoirs and three diversions to power and irrigation. As noted previously, the target diversion at all diversion points should be set equal to the element capacity. The results of the linear programming model yield target diversions less than element capacity due to limitations on total water, but for the purposes of the simulation

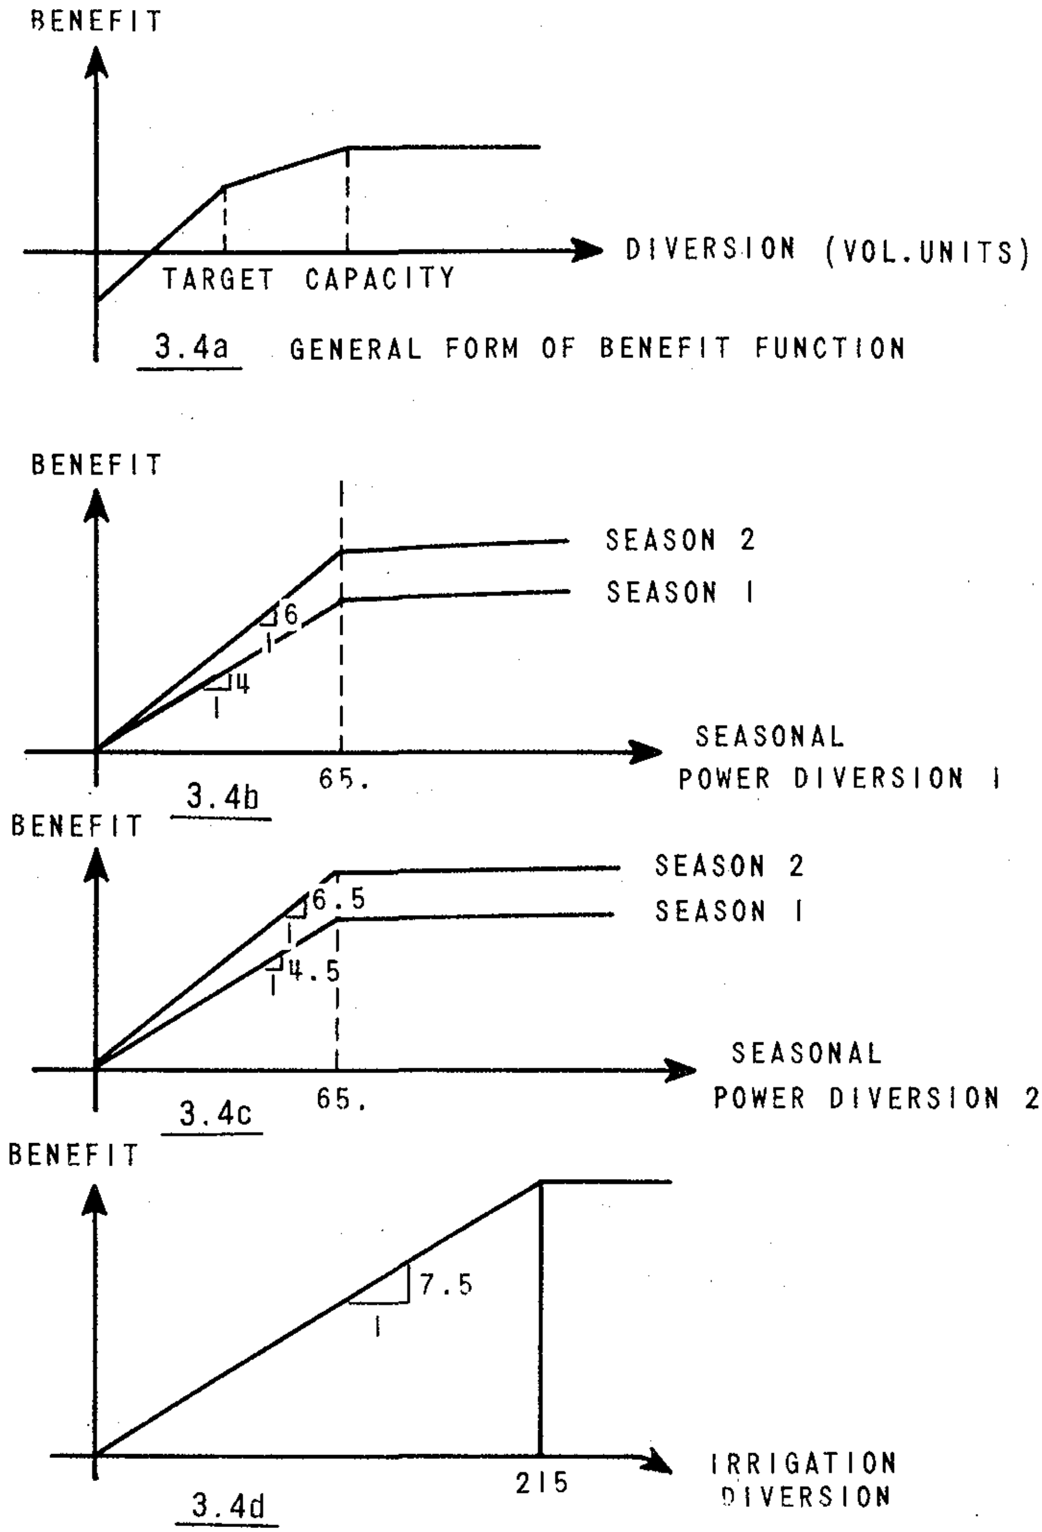

FIGURE 3.4

BENEFIT FUNCTIONS

TABLE 3.1

DATA FOR THE SIMULATION MODEL

		Season 1	Season 2
Mean Inflow (volume/period)			
	River 1	42.2	68.6
	River 2	39.0	62.4
Standard Deviation of Inflow (volume/period)			
	River 1	10.0	15.0
	River 2	10.0	15.0
Irrigation Requirement (vol./period)		175.0	40.0
Power Plant Capacity (volume/period)			
	Plant 1	65.0	65.0
	Plant 2	65.0	65.0

Reservoir	Maximum Capacity (volume units)
1	100.
2	100.
3	150.
4	200.

the targets are set to element capacity.

3.3 The Linear Programming Model

3.3.1 Introduction

As noted previously, results of Ibanez (15) point to the applicability of a linear programming model used in conjunction with a simulation. The main difficulty in constructing a linear programming model of an existing water resource system is generally in obtaining data and selecting the best possible linearization assumptions. In the present case, the simulation model was constructed to be as linear as possible. The only deviation from linearity in the simulation model is the stochastic nature of the inflows. Thus, the construction of the linear programming model was straight-forward. In particular, the use of only two hydrologic periods per year is a simplification of the linear programming formulation, since the resultant operation of the reservoirs is always to store in one season and to release in the other.

3.3.2 Formulation of the Linear Programming Model

The formulation of the linear programming model is in the terminology of the LP/MOSS system. The LP/MOSS system allows for the use of mnemonic names for the variables. For the present model, variables for each reservoir are the following:

1)	ViS, ViW (i=1,2,3,4)	Storage volumes in each of 4 reservoirs at end of Summer (1st period of year) and Winter seasons respectively
2)	TRGiS,TRGiW (i=1,2,3,4)	Summer and Winter releases at each reservoir
3)	QB1S,QB1W,QB2S,QB2W	Summer and Winter inflows to reservoirs 1 and 2 respectively

Similarly, variables for the power plant and irrigation facilities are:

1) PT1S,PT1W,PT2S,PT2W — Summer and Winter diversions to power Plants 1 and 2

2) CPS,CPW — Summer and Winter diversions to irrigation

In addition, for a complete formulation, the following variables are included:

1) FL1S,FL1W,FL2S,FL2W — Winter and Summer releases bypassing power plants 1 and 2

2) XSS,XSW — Summer and Winter excess water diverted from irrigation area

These last-defined variables are in reality slack variables defined to prevent negative flows at the power plants and irrigation area.

The above definitions are indicated schematically in Figure 3.5.

Two kinds of constraints are present in the model. The conservation constraints require conservation of flow at the reservoirs and diversions. The capacity constraints insure that capacities of reservoirs, power plants, and irrigation requirements are not exceeded.

Conservation constraints for the model are the following:

1)
$$\begin{aligned} \mathrm{V1S} &= \mathrm{V1W} + \mathrm{QB1S} - \mathrm{TRG1S} \\ \mathrm{V1W} &= \mathrm{V1S} + \mathrm{QB1W} - \mathrm{TRG1W} \end{aligned} \quad \text{Continuity at reservoir 1} \tag{3.1a}$$

2)
$$\begin{aligned} \mathrm{V2S} &= \mathrm{V2W} + \mathrm{QB2S} - \mathrm{TRG2S} \\ \mathrm{V2W} &= \mathrm{V2S} + \mathrm{QB2W} - \mathrm{TRG1W} \end{aligned} \quad \text{Continuity at reservoir 2} \tag{3.1b}$$

3)
$$\begin{aligned} \mathrm{V3S} &= \mathrm{V3W} + \mathrm{TRG1S} - \mathrm{TRG3S} \\ \mathrm{V3W} &= \mathrm{V3S} + \mathrm{TRG1W} - \mathrm{TRG3W} \end{aligned} \quad \text{Continuity at reservoir 3} \tag{3.1c}$$

4)
$$\begin{aligned} \mathrm{V4S} &= \mathrm{V4W} + \mathrm{TRG3S} + \mathrm{TRG2S} - \mathrm{TRG4S} \\ \mathrm{V4W} &= \mathrm{V4S} + \mathrm{TRG3W} + \mathrm{TRG2W} - \mathrm{TRG4W} \end{aligned} \quad \text{Continuity at reservoir 4} \tag{3.1d}$$

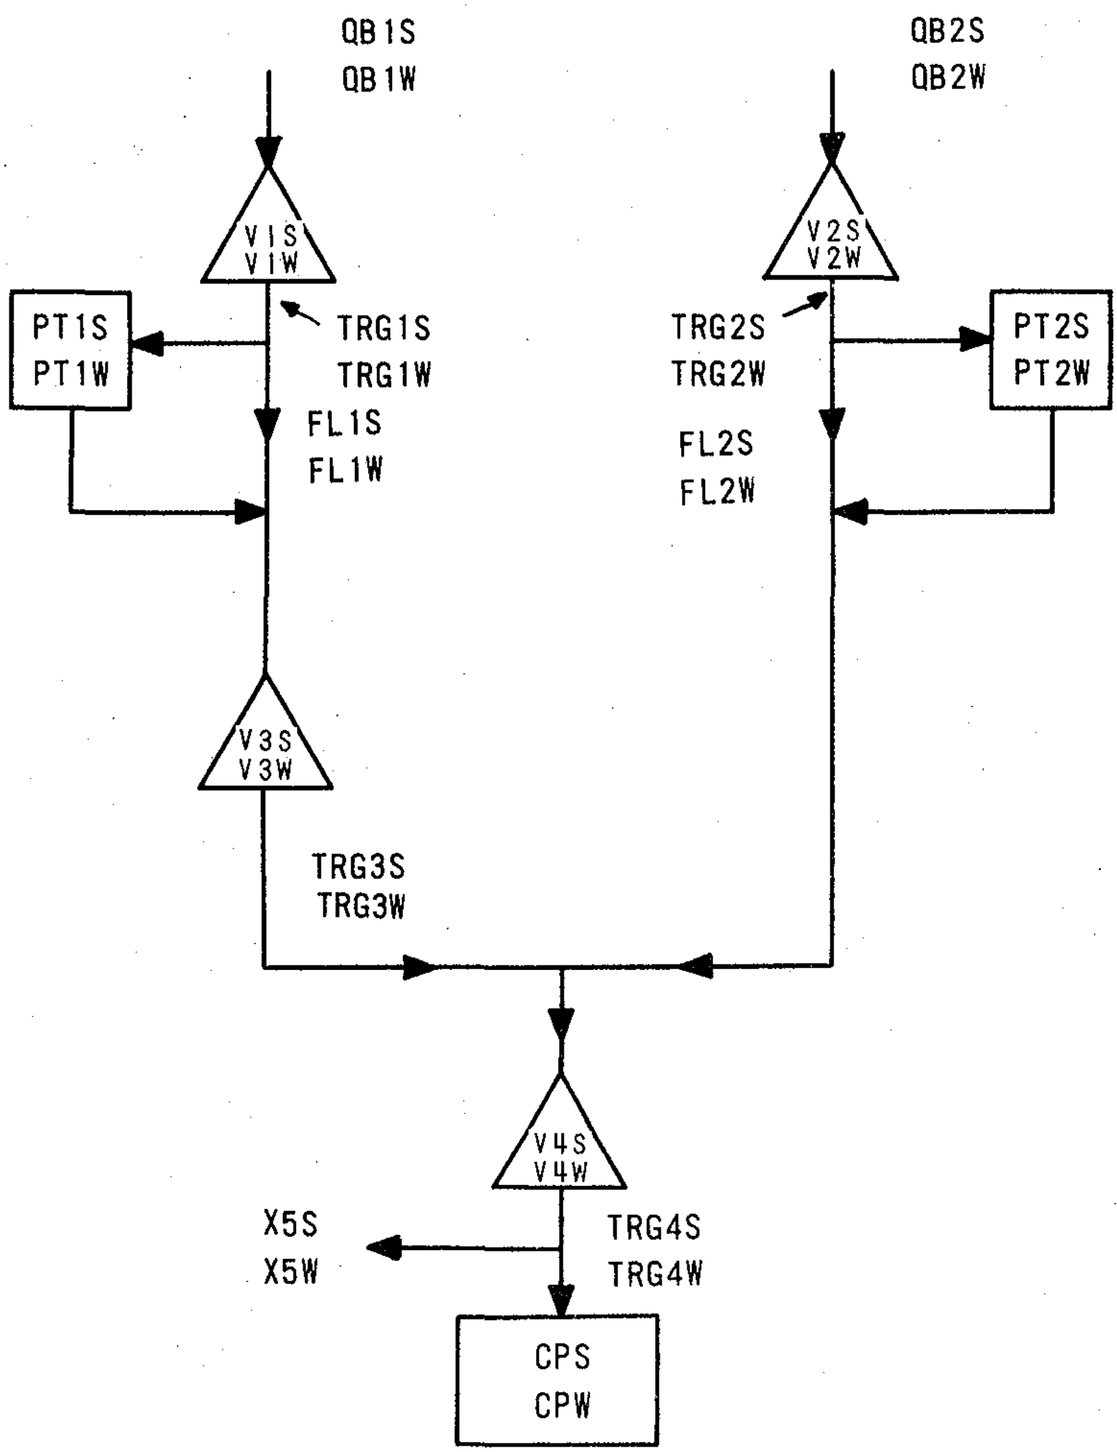

FIGURE 3.5

LINEAR PROGRAMMING MODEL

5) TRG1S = PT1S + FL1S

TRG1W = PT1W + FL1W

Continuity at Plant 1 (3.1e)

6) TRG2S = PT2S + FL2S

TRG2W = PT2W + FL2W

Continuity for Plant 2 (3.1f)

7) TRG4S = CPS + XSS

TRG4W = CPW + XSW

Continuity for Irrigation (3.1g)

The above constraints are sufficient to define the configuration of the linear programming model, and the scheme of water transfers. Capacity constraints for the model are the following:

1)	V1S, V1W	$\leq$	Capacity of Reservoir 1	(3.2a)
2)	V2S, V2W	$\leq$	Capacity of Reservoir 2	(3.2b)
3)	V3S, V3W	$\leq$	Capacity of Reservoir 3	(3.2c)
4)	V4S, V4W	$\leq$	Capacity of Reservoir 4	(3.2d)
5)	PT1S, PT1W	$\leq$	Capacity of Plant 1	(3.2e)
6)	PT2S, PT2W	$\leq$	Capacity of Plant 2	(3.2f)
7)	CPS	$\leq$	Summer Irrigation Demand	(3.2g)
8)	CPW	$\leq$	Winter Irrigation Demand	(3.2h)

Constraints 1 through 4 insure that storage will not exceed reservoir capacity in any reservoir in either summer or winter seasons. The remainder of the capacity constraints insure that the power and irrigation capacities are not exceeded.

Equations (3.1) and (3.2) are consistent with a discretization procedure for the model which corresponds to that used in the simulation. Inflows may be considered to take place at the beginning of a period, and drafts released immediately thereafter. The reservoirs then store at steady state for the remainder of the period. The draft may have any value up to the sum of the inflow plus stored volume from the previous period, as an examination of equations (3.1) will show. This discretization procedure is equivalent to the assumption of uniform draft and inflow rates throughout the period.

The objective function, given the label PROFIT, is taken as:

$$\text{PROFIT} = 4.0\ \text{PT1S} + 6.0\ \text{PT1W} + 4.5\ \text{PT2S} + 6.5\ \text{PT2W} + 7.5\ (\text{CPS}+\text{CPW}) \quad (3.3)$$

The value of the coefficients is identical to that used in the benefit functions in the simulation model, but no discounting is present. As in the simulation, the benefit is solely a function of the diversions to irrigation and power. The objective function is to be maximized.

In running the linear programming model, QB1S, QB1W, QB2S, and QB2W are constrained to be equal to the seasonal means obtained from the 5-year hydrology originally selected for the simulation, given in Table 3.1, and the reservoir, power plant and irrigation capacities are constrained to the values shown in Table 3.1. All variables are taken as non-negative.

3.3.3 Results of the Linear Programming Model

Implementation of the linear programming model on the IBM 1130 required 18 minutes, start-to-finish, and required 16 Simplex iterations. The maximum value of the objective function was 2751 benefit units. The values of the other variables are shown schematically in Figure 3.6, and given in tabular form in Table 3.2. Results show that the summer irrigation capacity and winter power capacity are used in full, while the winter irrigation and summer power are at less than their maximum values, due to insufficient water. Reservoirs 1, 3, and 4 store in the winter, and reservoir 2 stores in the summer. There is no excess flow, as evidenced by the zero values of the slack variables. The capacity constraints on the reservoirs were not binding, and the results indicate reservoirs much smaller than the capacity constraints. The capacities used in the following studies are both the capacity constraint values, and later the linear programming capacities.

3.4 Studies with the Simulation Model

3.4.1 *Scope of the Studies*

As noted in the introduction, the primary investigation to be performed is the study of the proposed LPSS technique. Additional studies examine in a limited manner the characteristics of a simulation model with an associated linear programming model, and test features which cannot be readily examined with the larger models.

Studies have been grouped into different series, each series comprising one major area of investigation. For all but one of the studies, simulation runs of 50 years were used, and only summary output was obtained. Benefit functions and power and irrigation capacities were held constant for all studies, while initial volume, interest rate, capacities, and operating policy targets were taken as parameters. The parameters of the operating policy were chosen as the 8 target drafts, 1 for each of the 4 reservoirs for each period, yielding an 8-dimensional policy parameter vector. Each simulation run of 50 years required approximately .02 minutes of computation time on the IBM 360/65 computer. Thus, a large number of runs could be made without prohibitive cost in terms of computer time. Descriptions of the type of study made in each series are presented in the remainder of section 3.4, results and discussion in section 3.5, and conclusions for this model are found in section 3.6.

3.4.2 Series I

Series I studies compare results for initial parameter sets obtained

TABLE 3.2

RESULTS OF LINEAR PROGRAMMING MODEL

Variable	Value (volume units)
V1S	0.0
V1W	3.6
V2S	2.6
V2W	0.0
V3S	0.0
V3W	65.0
V4S	0.0
V4W	27.8
TRG1S	45.8
TRG1W	65.0
TRG2S	36.4
TRG2W	65.0
TRG3S	110.8
TRG3W	0.0
TRG4S	175.0
TRG4W	37.2
PT1S	45.8
PT1W	65.0
PT2S	36.4
PT2W	65.0
CPS	175.0
CPW	37.2
FL1S	0.0
FL2S	0.0
FL1W	0.0
FL2W	0.0
XSS	0.0
XSW	0.0

PROFIT = 2751 benefit units

Reservoir	Capacity (volume units)
1	3.6
2	2.6
3	65.0
4	27.8

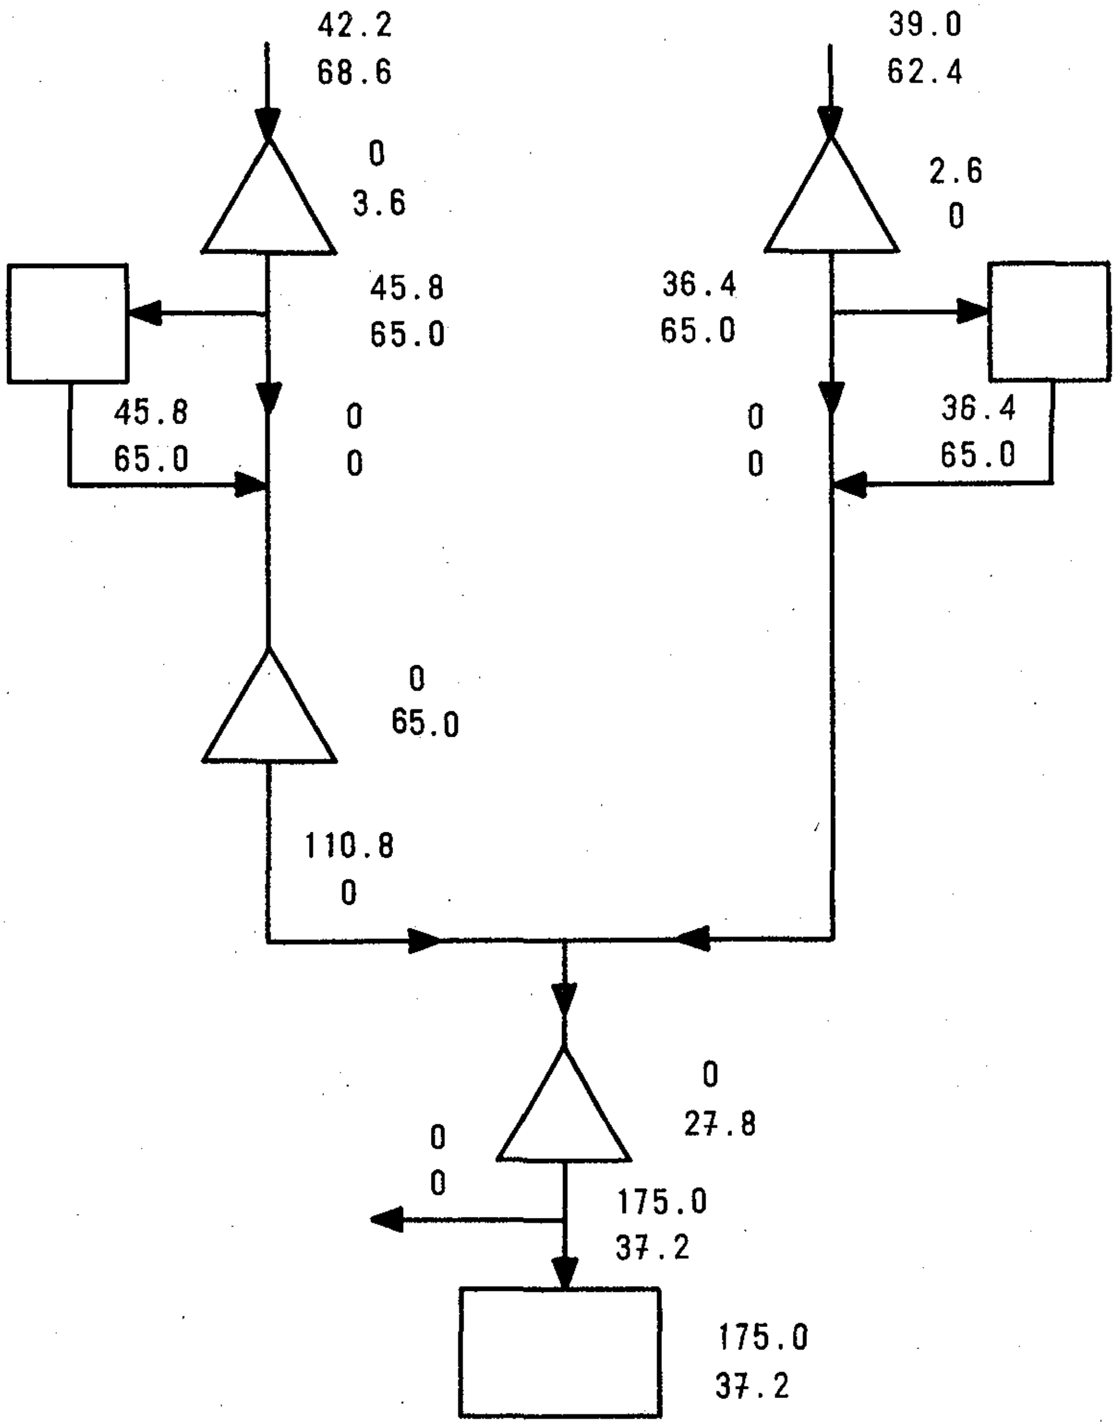

FIGURE 3.6
LP SOLUTION FOR 2--PERIOD MODEL

by linear programming and random sampling. Influence of interest rate and initial volume is studied briefly to select appropriate values for use in the remaining portion of the LPSS process.

Reservoir capacities were set at 100, 100, 150 and 200 volume units for reservoirs 1 through 4 respectively, using data set 2. A total of seven parameter sets were studied with a view to selecting a starting point for hill-climbing. The values of the 8 individual targets for each parameter set are presented in Table 3.3a. (The various portions of Table 3.3 present values of the parameter sets used in the studies, and hereafter, parameter sets will be referred to by their numbering in Table 3.3).

Parameter sets 1, 2, and 3 define policies with targets selected at random from the range of "reasonable" values for each of the targets. Set 4 is a policy at the origin in parameter spcae, and operates to maintain full reservoirs by setting controlled release at zero, requiring all releases to be spills. Set 5 is the inverse of this, and causes operation with empty reservoirs, since the value of 200 units for target draft is always greater than the available water for any period. If initial reservoir volume is zero, parameter set 5 yields a limiting policy of run of the river operation. After the reservoirs are either full or empty, operation according to sets 4 and 5 is identical. Set 6 is the parameter set derived from the linear programming model. Set 7 was obtained after examining the results of simulation runs with sets 1 through 6, and as such is not strictly a "starting value", since some degree of search has been performed to obtain set 7. Sets 1 and 6 yielded highest benefits. Parameter set 7 was constructed by starting from set 6 and "moving in the direction of"

TABLE 3.3

OPERATING POLICY PARAMETER SETS

Policy	T(1,1)	T(2,1)	T(3,1)	T(4,1)	T(1,2)	T(2,2)	T(3,2)	T(4,2)
1	90	80	100	180	80	60	60	60
2	70	30	85	105	30	40	65	65
3	60	50	70	150	50	30	30	30
4	0	0	0	0	0	0	0	0
5	200	200	200	200	200	200	200	200
6	45.8	36.4	110.8	175	65	65	0	37.2
7	65	65	110	175	65	65	0	40

Table 3.3a

Parameter sets for Initial Studies

Policy	T(1,1)	T(2,1)	T(3,1)	T(4,1)	T(1,2)	T(2,2)	T(3,2)	T(4,2)
8	50.8	36.4	110.8	175	65	65	0	32.2
9								37.2
10								38.0
11								39.0
12								40.0
13								41.0
14								42.2

Table 3.3b

Parameter Sets for Sensitivity to Target (4,2)

TABLE 3.3 (CONT.)

Policy	T(1,1)	T(2,1)	T(3,1)	T(4,1)	T(1,2)	T(2,2)	T(3,2)	T(4,2)
15	65	65	110	175	65	65	0	40
16	70							
17	60							
18	65	70						
19		60						
20		65	115					
21			105					
22			110	180				
23				170				
24				175	70			
25					60			
26					65	70		
27						60		
28						65	5	
29							0	45
30								35

Table 3.3c

Parameter Sets for Single-Factor analysis

TABLE 3.3 (CONT.)

Policy	T(1,1)	T(2,1)	T(3,1)	T(4,1)	T(1,2)	T(2,2)	T(3,2)	T(4,2)
31	50	35	*	175	65	65	0	40
32	55							
33	60							
34	65							
35	50	40						
36	55							
37	60							
38	65							
39	50	45						
40	55							
41	60							
42	65							
43	50	50						
44	55							
45	60							
46	65							
47	50	55						
48	55							
49	60							
50	65							
51	50	60						
52	55							
53	60							
54	65							
55	50	65						
56	55							
57	60							
58	65							

* - The values of target (3,1) depends on results of the single-factor and sensitivity analysis in each case.

Table 3.3d

Parameter Sets for Marginal Analysis

TABLE 3.3 (CONT.)

Policy	T(1,1)	T(2,1)	T(3,1)	T(4,1)	T(1,2)	T(2,2)	T(3,2)	T(4,2)
59	65	65	112	175	65	65	0	40
60			114					
61			116					
62			118					
63			120					
64			130					
65			140					
66			150					
67			160					

Table 3.3e

Parameter Sets for Sensitivity to Target (3,1)

Set 1 in policy parameter space.

3.4.3 Series II

Series II constitutes the continuation of the LPSS technique from a starting point located by Series I. Reservoir capacities are those used in Series I, with an interest rate of 0. Initial volumes are taken as the appropriate steady-state contents of each reservoir, as determined by the linear programming model.

Six synthetic hydrologies were generated and used. Set 7 was used as the starting point on the basis of results of Series I, and a single-factor and marginal analysis were performed for three of the hydrologies.

3.4.4 Series III

Series III studies are similar to those of Series II, except that capacities used in the simulation are those generated by the linear programming model, presented in Table 3.2. A starting set for the search is again selected as set 7, and results of single-factor, marginal, and axial analysis are presented in graphic and tabular form. Synthetic hydrologies are used, and the LPSS process carried out in detail for hydrologies 1 and 2. A portion of the Series III study examines the problem of differing "near-optimal" parameter sets being obtained for different hydrologies by the LPSS process.

3.4.5 Series IV

The foregoing studies served to test the general concept of LPSS. The remaining series examine sensitivity of the LPSS process to initial volume, capacity, and interest rate, and study interfacing of the linear programming and simulation models.

Series IV investigates the sensitivity of the near-optimal parameter set obtained as the result of an LPSS process to changes in the reservoir capacity. If a near-optimal set determined for LP capacities can be shown insensitive to small changes in reservoir capacities, then the planning effort is reduced considerably, since an LPSS process need not be performed for each small change in proposed reservoir capacity.

Series IV studies were undertaken with hydrology 1 (of the six synthetic hydrologies) and initially empty reservoirs, for an interest rate of 0. To reduce the number of parameters, variation in reservoir capacity was obtained by multiplying the linear programming capacities by a single parameter, thus maintaining the ratio of reservoir capacities identical for all runs, and allowing plotting of results as a function of the single multiplying factor.

3.4.6 Series V

Studies were undertaken in Series V to investigate sensitivity of the near-optimal parameter set obtained by LPSS for a 50-year simulation with linear programming initial volumes, to changes in initial volume. Hydrology 1 was used for linear programming reservoir capacities and various interest rates. As in Series IV, the number of parameters was reduced by taking initial volumes as a single percentage of reservoir capacity for all reservoirs. This percentage is the parameter against which benefits are plotted.

3.4.7 Series VI

Runs made in Series VI are an investigation of the sensitivity of the near-optimal policy obtained by LPSS for an interest rate of 0.0, to changes in interest rate. Various interest rates are studied. For each of the interest rates, an LPSS process is carried out to determine the near-optimal policy for that particular interest rate. Results are compared to the results of simulation runs using the near-optimal policy found for 0.0 interest rate at the different interest rate. Hydrology 1 is used, with initially full reservoirs and linear programming reservoir capacities.

3.4.8 Series VII

Series VII studies examine the persistence of the effects of changes in initial volume, using the near-optimal policy generated by LPSS for the full 50-year simulation with LP initial volumes. Hydrology 1 is investigated at an interest rate of 0.0. Benefits for the cases of initially full and initially empty reservoirs are compared for various durations of simulation. Together with the results of Series V, Series VII studies serve to investigate the possibility of developing state-orientation in the operating rule through the simple expedient of variation in initial volume.

3.4.9 Series VIII

A suggested technique for testing the correlation between the linear programming and simulation models of a basin is examined in the study of Series VIII. The proposed technique is to examine the benefits in each model for zero reservoir capacity, thus entirely eliminating operation as a factor, and isloating the effects of stochastic inflow and non-linearities. A run is made with zero reservoir capacity for 50 years duration, and the mean annual benefit is compared with that derived from the linear programming model with run of the river operation.

3.5 Results and Discussion of Results

3.5.1 Series I

Results of Series I studies to compare initial parameter sets from random sampling and linear programming are presented in Table 3.4. For the studies of Series I only, the summer power capacity for parameter set 6 was changed from 65 units for both plants to the values determined by linear programming, 45.8 units for plant 1 and 36.4 units for plant 2. The reduced power target values are generated by linear programming in accordance with constraints on the available water. Thus, in each summer period in the steady-state linear programming solution, the amount of water at the diversion point is limited by the hydrology to the values of 45.8 and 36.4 units. Since no more water is available, power targets are set by linear programming to this value. For the simulation however, stochasticity in the inflows occassionally produces flows at the diversion points in excess of the values prescribed by linear programming. To make use of this additional water, the power targets in the simulation model are ordinarily set to 65 units, equal to the plant capacities. For Series I only, the precise results of the linear programming model were used to establish the power targets. A test was performed to determine the decrease in benefits resulting from using the LP-generated power targets. Results indicate a decrease in benefits of 128 units from benefits produced using the 65 unit targets throughout, for an interest rate of 0.

Figure 3.7 plots a portion of the data of Table 3.4, as benefit vs. interest rate for parameter set 7 and LP initial volumes. This plot shows the typical form for such a discounted evaluation, and indicates the extreme effect of interest rate in determining the level of benefits. An examination of Table 3.4 shows that the rank ordering of the 7 parameter sets according to associated benefits does not change for the two different interest rates of 0.0 and 0.06.

Figure 3.8 is a comparison of the effects of initial volume on the 7 parameter sets studied, for an interest rate of .06. Benefit is plotted

POLICY	INT.	INITIAL VOLUME = ½ CAPACITY	INITIAL VOLUME = L.P. VALUE
1	.06	23048	22405
2		18601	18185
3		21273	20697
4		18928	18245
5		15928	14811
6		23670	23426
7		24240	23593
1	.00	—	127433
2		—	108985
3		—	124377
4		—	101589
5		—	97845
6		137326	136230
7		138313	136674
6	.05	—	27792
7		—	27569
6	.04	—	34090
7		—	34288
6	.03	—	43693
7		—	43924
6	.02	—	59266
7		—	59550
6	.01	—	86250
7		—	86609

TABLE 3.4

RESULTS OF SERIES I

BENEFIT FOR DIFFERENT POLICIES

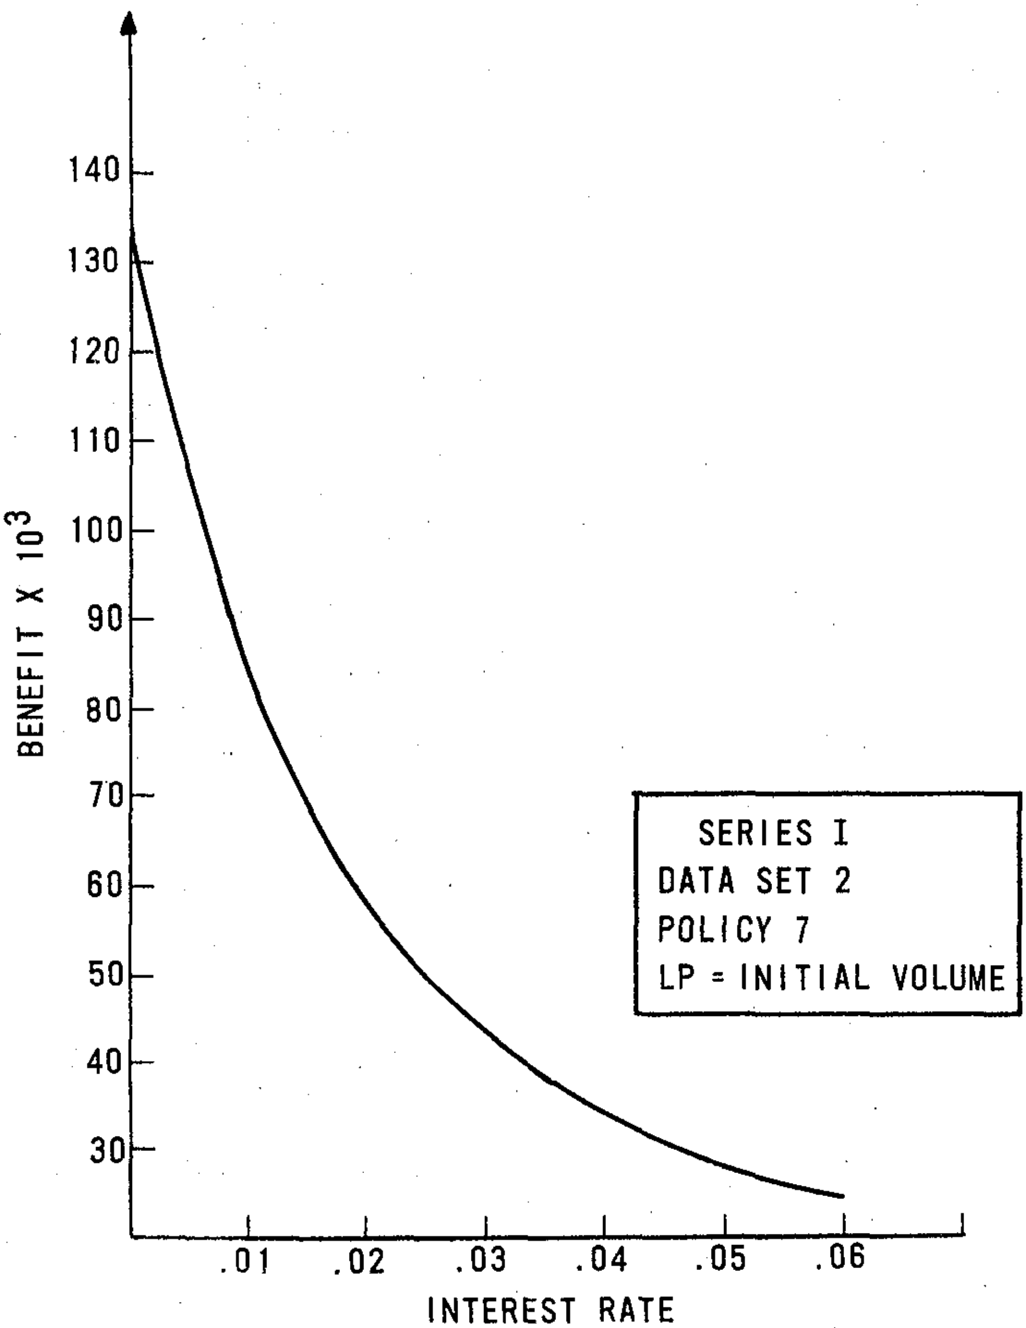

FIGURE 3.7

BENEFIT VS. INTEREST RATE

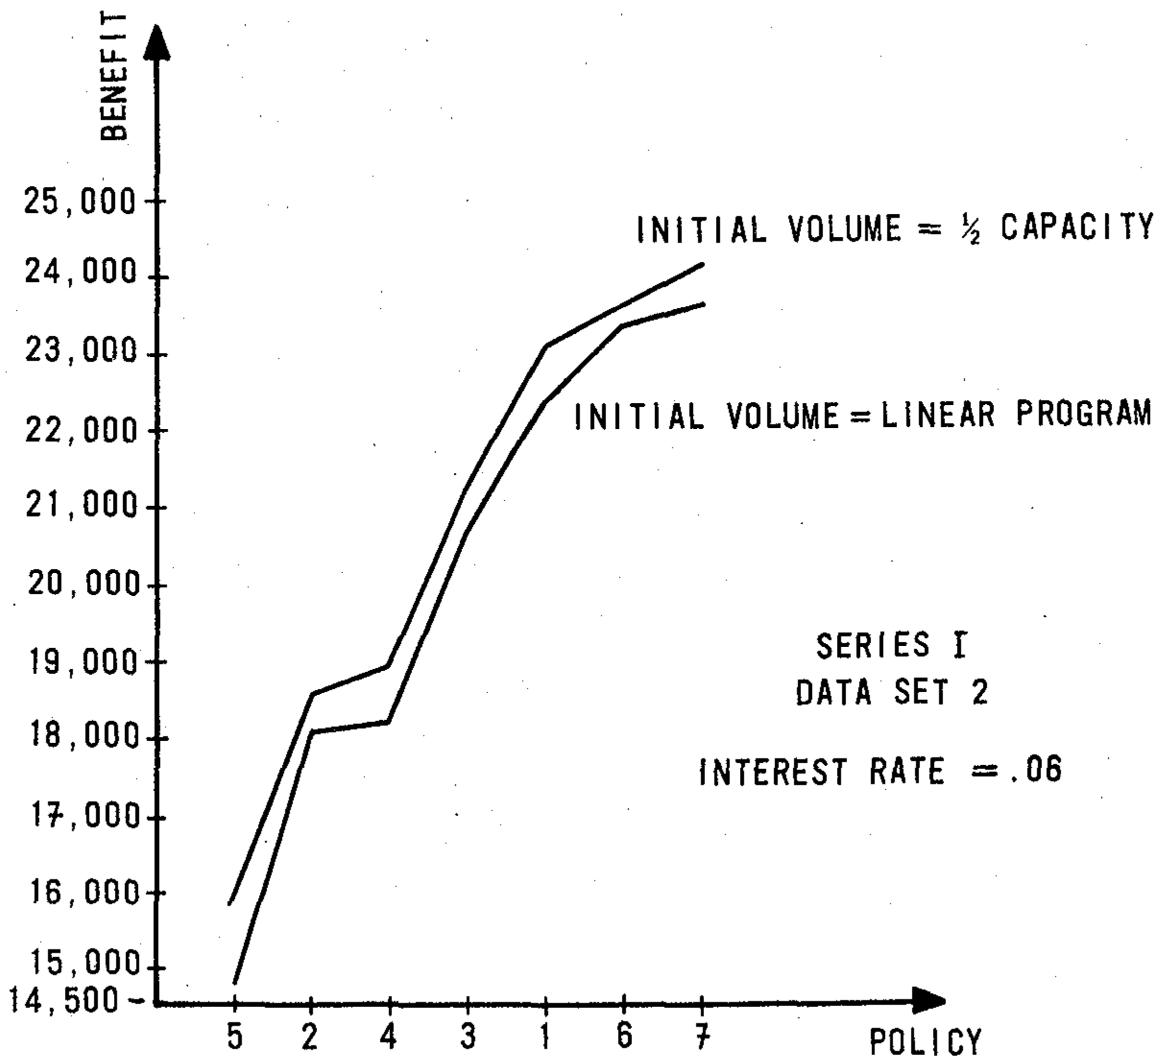

FIGURE 3.8

BENEFITS VS POLICY , INITIAL VOLUME AS PARAMETER against "policy" with inital volume as a parameter. The "policy" axis places parameter sets in rank order for decreasing benefits. As can be seen from the figure, no change in rank ordering is found for the two different values of initial volume, and no significant difference in the shape of the curves is noted.

Results of Series I show that initial values for an LPSS process may be otained with a simple procedure. After completion of the simulation model, formulation of the linear programming model provee quite straightforward. Of the 6 parameter sets initially tested, linear programming yielded the highest benefits, but the presence of the random policies served to indicate areas of improvement over the linear programming parameter set, which led to parameter set 7, the principal starting point for further search. Early identification of the initial volume and interest rate as parameters allowed for selection of appropriate values for later studies.

The fact that rank-ordering of different parameter sets did not change for two different interest rates led to the selection of 0.0 as an appropriate value for the majority of the LPSS studies. An interest rate of 0,0 leads to highest benefits, and consequently, greatest resolution of the difference between methods of operation. In addition, it allows comparisons of mean annual benefits for simulation runs with the annual benefit obtained by linear programming. Thus, mean annual benefits for a 50-year simulation run with parameter set 6 of 2724 units are very close to the annual benefit of 2751 units obtained by the linear programming model. The large reservoir capacities serve to regulate the flow and reduce the effects of stochasticity so that the two models are comparable. The manner of construction of data set 2 from 10 replicate sets of 5-year hydrologic sequences further reduces the effect of stochasticity that would be found in a sequentially generated synthetic hydrology.

In the absence of information about initial volumes, those values indicated by the linear programming solution are used for the majority of the LPSS studies. Figure 3.8 shows that the rank ordering of the parameter sets is insensitive to initial volume over a 50-year simulation, so this decision does not appear to be critical. If the situation is such that a particular value of initial volume is known or expected to exist, it should be used in place of the linear programming volume.

3.5.2 Series II

Results of Series II studies of the LPSS process for 6 synthetic hydrologies are presented in Table 3.5. Table 3.5a presents benefits obtained for a 50-year simulation run for the 7 initial parameter sets for 6 different hydrologies, with an interest rate of 0.0. Rank ordering of the 7 parameter sets with respect to benefits does not change as a function of hydrology used. Parameter set 7 remains the best initial point.

Studies of Series I and II suggested that the value of target (4,2) could be fixed for all parameter sets by a sensitivity analysis. Parameter sets 8 through 14, presented in Table 3.3b, were used to obtain the results presented in Figure 3.9, for benefits as a function of target (4,2). These results indicate clearly that the value of target (4,2) should be fixed at 40 units, equal to the 2nd period irrigation capacity. Physical reasoning confirms the correctness of this value.

Table 3.5b presents results of a single-factor analysis from the starting point of parameter 7. This single-factor analysis constitutes the second step in the LPSS process. Parameter sets refer to those given

TABLE 3.5

interest = 0.0 LP initial volume Reservoir Capacities 100,100,150,200

RESULTS OF SERIES II

Policy	1	2	3	4	5	6 ←(Hydrology)
1	127173	125866	129023	124359	129156	128534
2	108370	107684	109109	105968	109697	109145
3	123805	123417	124934	122354	125162	124526
4	96220	95372	97431	94721	98738	97032
5	99832	98768	101192	98420	102340	101101
6	135335	133864	136804	133364	135975	135679
7	136281	134240	137427	133524	137178	137135

Table 3.5a

Benefits for Initial Studies, 6 Hydrologies

policy	1	2 ←(hydrology)
15	136281	134240
16	136196	134105
17	136243	134236
18	136141	134131
19	136317	134260
20	136580	134478
21	136114	134203
22	135347	133582
23	135328	134091
24	135784	133766
25	135850	133830
26	135805	133807
27	136021	134010
28	136580	134478
29	134506	132688
30	135291	134053

Table 3.5b

Single Factor Analysis, 2 Hydrologies

TABLE 3.5 (CONT.)

Policy	1	2	3 ←(hydrology)
31	135817	134100	137063
32	135934	134582	137154
33	135893	134624	137086
34	135909	134628	137124
35	136390	134118	138002
36	136507	134570	138027
37	136466	134611	137926
38	136525	134615	137909
39	136456	133987	137828
40	136573	134438	137853
41	136532	134480	137752
42	136571	134484	137755
43	136376	133840	137678
44	136493	134291	137703
45	136452	134333	137602
46	136511	134337	137605
47	136285	133788	137591
48	136402	134239	137616
49	136361	134280	137514
50	136419	134285	137518
51	136210	133769	137537
52	136327	134220	137562
53	136286	134261	137460
54	136345	134266	137464
55	136173	133749	137512
56	136291	134201	137531
57	136249	134242	137436
58	136308	134246	137439

Target (3,1) = 110.8 for all hydrologies

Table 3.5c

Marginal Analysis, 3 hydrologies

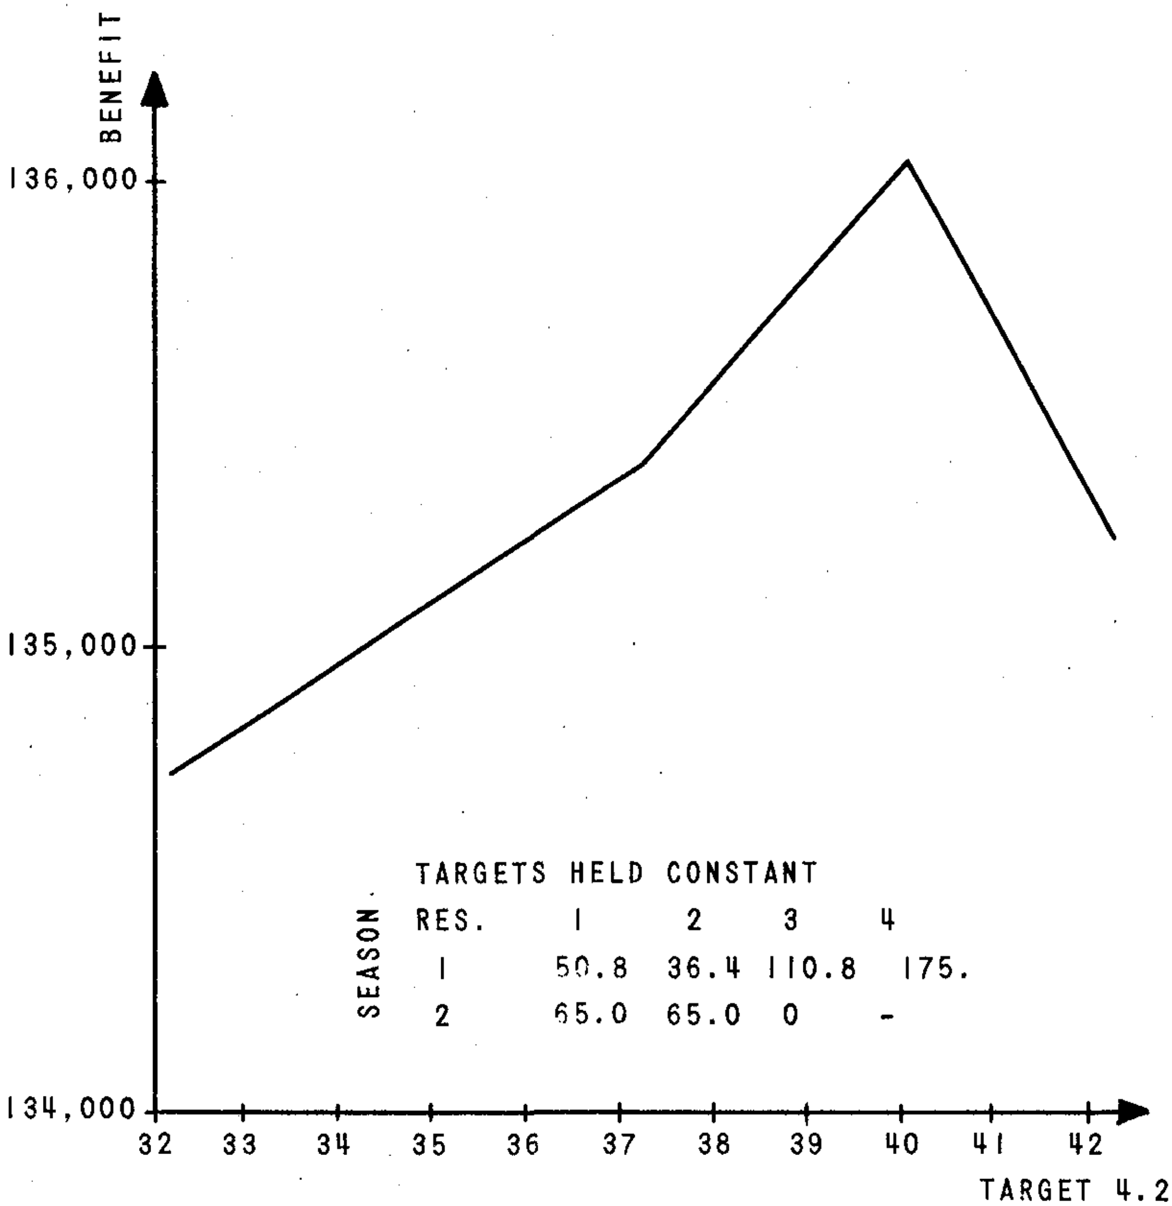

HYDROLOGY 1 INTEREST RATE:0. LP INITIAL VOLUMES

FIGURE 3.9

<u>SENSITIVITY OF BENEFITS TO TARGET (4.2)</u>

in Table 3.3c. Results show that for both hydrologies 1 and 2, greatest benefits are obtained by parameter sets 20 and 28, which represent increases of 5 units in targets (3,1) and (3,2) respectively. Benefits obtained are identical since the two parameter sets both prescribe the release of an additional 5 units of water for irrigation, and the downstream regulation afforded by reservoir 4 provides that the water supplied in the 2nd season can be stored for use in the first season. Thus, although the physical operation of the basin due to parameter sets 20 and 28 is different, the net effect of both is to supply an additional 5 units of irrigation water in the first period, and hence benefits are identical. Since the net effect is to provide additional water in the first season, further studies are carried out with target (3,2) at zero level.

Table 3.5c presents results of a marginal analysis on targets (1,1) and (2,1). These targets were selected as being significant after examination of the results of the single-factor analysis. For the marginal analysis, the value of target (3,1) was held constant at 110.8 units.

Marginal analysis yields different combinations of the two targets for each of the three hydrologies. Thus, the targets (1,1) and (2,1) for hydrology 1 are 55 and 45 units, those for hydrology 2 are 65 and 35 units, and those for hydrology 3 are 55 and 40 units.

The above procedure does not represent a true LPSS search, since the value of target (3,1) in the marginal analysis was not selected by sensitivity analysis for each hydrology. The true LPSS search was not performed due to the recognition that the reservoir capacities for Series II are excessive. Accordingly, Series III presents results of a true LPSS search with appropriate reservoir capacities.

Notwithstanding the above consideration, certain features of the LPSS process can be noted. Levels of benefit obtained for 50-year simulation runs are not significantly affected by using different hydrologies. Comparison of results for the first 7 parameter sets with results presented in Table 3.4, which used data set 2 as the hydrology specification, shows that benefits are again comparable, even though the hydrologies were generated by different means.

Although benefits are comparable for different hydrologies, results of the marginal analysis show that the near-optimal parameters are a function of hydrology. This dependence is examined further in Series III.

3.5.3 Series III

Series III studies constitute a full LPSS process. Reservoir capacities generated by linear programming were used, instead of the large capacities of Series I and II. As a consequence, stochastic effects are more significant, and the reservoirs will spill from time to time, which was rarely the case in the simulation runs of Series I and II.

Table 3.6a presents results of simulation runs for 6 hydrologies for the 7 initial parameter sets. Behavior is similar to that found in Series I and II, with rank ordering being constant for all hydrologies, and parameter set 7 being superior in all cases,

Single-factor analysis, with results presented in Table 3.6b, shows that benefits are most sensitive to target (3,1). Contrary to results of Series II, increases in target (3,2) cause decreases in benefits. The difference is due to the decreased capacity for regulation by reservoir 4 in Series III, favoring targets which supply irrigation water when it is needed, in period 1.

TABLE 3.6

RESULTS OF SERIES III

Policy	1	2	3	4	5	6	←(hydrology)
1	115086	114098	116673	112496	116830	116464	
2	109380	108315	109685	107282	110209	109786	
3	121238	120105	122074	119903	122721	121789	
4	99260	98148	100520	97851	101704	100333	
5	99832	98768	101182	------	102340	101101	
6	130077	129777	131747	------	131875	131238	
7	130416	130119	131878	------	132151	131669	

Table 3.6a

Initial Studies, Linear Programming Capacities, 6 hydrologies

policy	1	2	←(hydrology)
15	130416	130119	
16	130416	130119	
17	130416	130114	
18	130385	130119	
19	130399	130119	
20	130455	130311	
21	130239	129730	
22	130250	130039	
23	130029	129633	
24	130269	130003	
25	130402	130072	
26	130274	129966	
27	130405	130148	
28	130083	129643	
29	129736	129329	
30	130378	130038	

Table 3.6b

Single Factor Analysis, 2 hydrologies

TABLE 3.6 (CONT.)

policy	1	2 ←(hydrology)
59	130459	130207
60	130461	130281
61	130455	130341
62	130454	130390
63	130437	130431
64	130437	130541
65	130437	130569
66	130437	130569
67	130437	130569

Table 3.6c

Sensitivity to Target (3,1), 2 hydrologies

policy	1*	2** ←(hydrology)
31	130186	130311
32	130242	130339
33	130241	130338
34	130241	130344
35	130245	130354
36	130300	130383
37	130300	130382
38	130300	130387
39	130330	130408
40	130385	1304[illegible]7
41	130385	130446
42	130385	130452
43	130382	130485
44	130437	130531
45	130437	130539
46	130437	130545
47	130413	130516
48	130468	130561
49	130468	130577
50	130468	130582
51	130412	130503
52	130468	130548
53	130467	130564
54	130467	130569
55	130406	130508
56	130462	130548
57	130461	130564
58	130461	130569

* - target (3,1) = 114

** - target (3,1) = 140

Table 3.6d

Marginal Analysis, 2 hydrologies

The next step in the LPSS process was to perform a sensitivity analysis to values of target (3,1), for hydrologies 1 and 2. Results are presented in Table 3.6c, with parameter sets being identified by corresponding numbers in Table 3.3e. For hydrology 1, results indicate a value of 114 units, while for hydrology 2 140 units yields maximal benefits.

After tentative values for target (3,1) have been fixed for both hydrologies, the marginal analysis of Table 3.6d is performed. Results are plotted in Figures 3.10 and 3.11 for hydrologies 1 and 2 respectively, as plots of benefit vs. target (2,1), with target (1,1) as a parameter.

Figure 3.10 shows a fairly large flat region of benefit response for target (1,1) and target (2,1) in the range of 55 to 65 units. Identification of such regions of indifference is of great importance since factors other than economic may be called into play to determine precise target values within the indifference region. Further, an indifference region in the response for a particular hydrology is helpful in selecting a single near-optimal parameter set for various hydrologies. This is pointed out by a comparison of Figures 3.10 and 3.11. Figure 3.11 shows a maximum at target (1,1) of 65 and target (2,1) of 55 units. The flatness of the response in Figure 3.10 also allows the selection of these values for the near-optimal parameter set. The only remaining difference between the parameter sets for hydrologies 1 and 2 is reduced to the difference in value of the single variable target (3,1).

In light of the above, sensitivity of the near-optimal parameter set to values of target (3,1) was studied for both hydrologies, with results presented in Figure 3.12. Such a study serves two purposes.

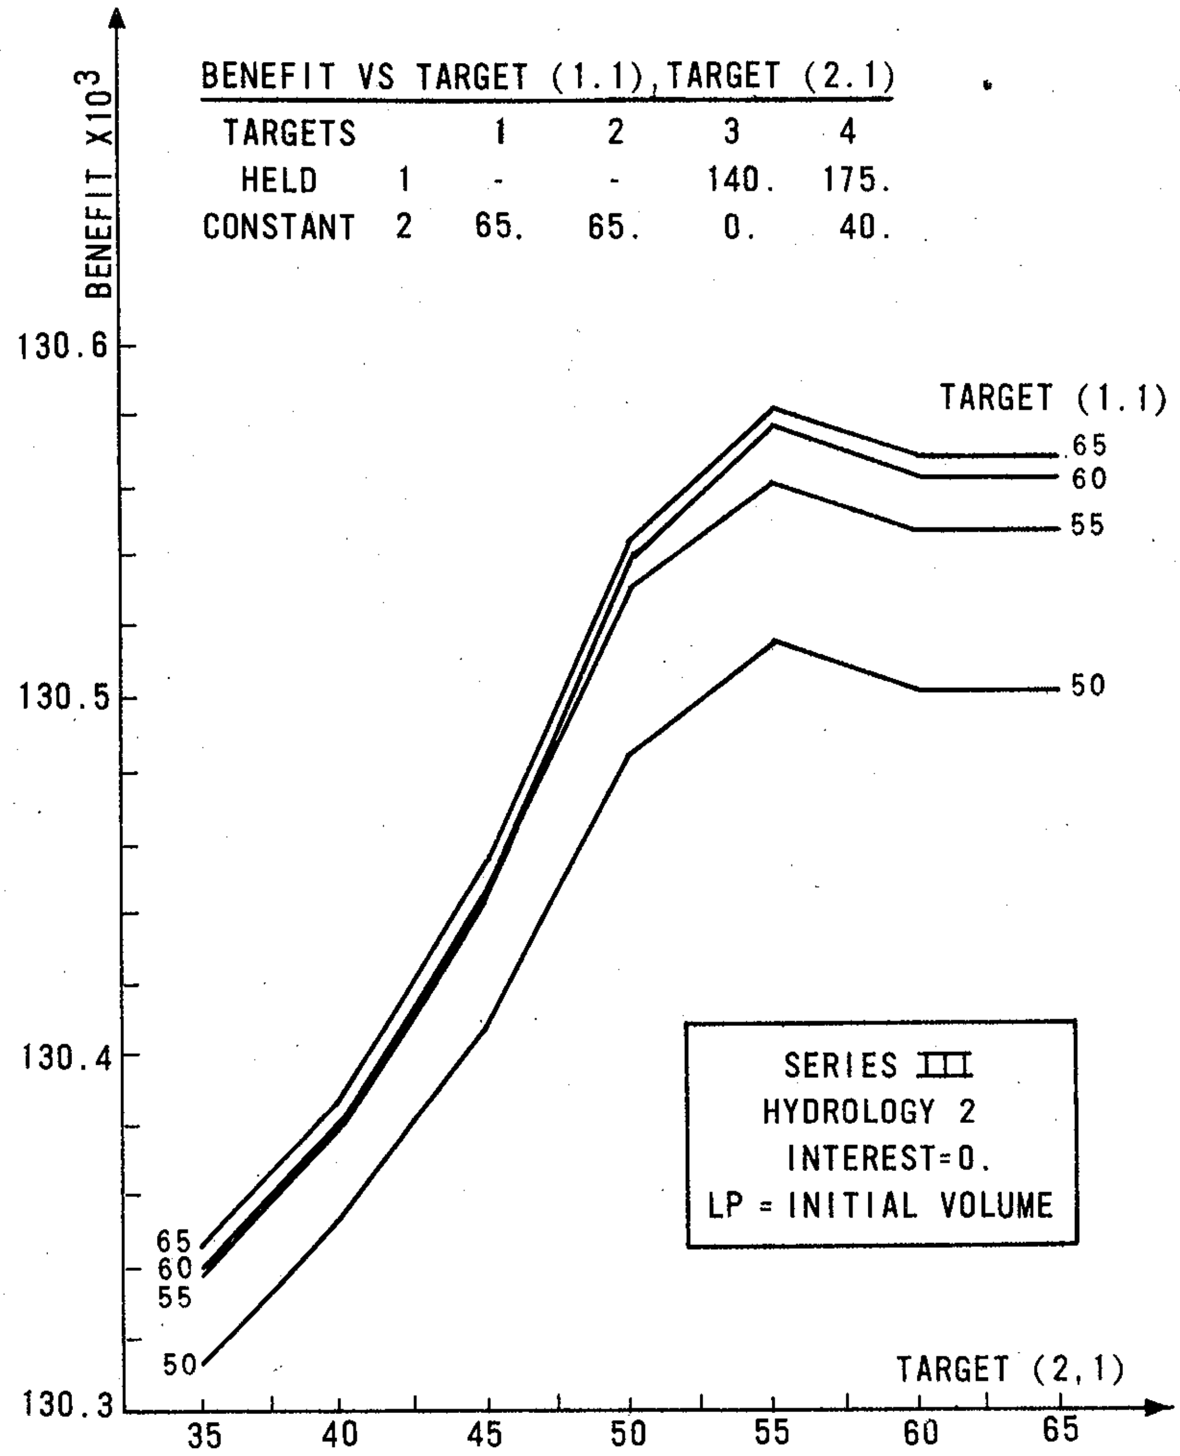

FIGURE 3.10

MARGINAL ANALYSIS

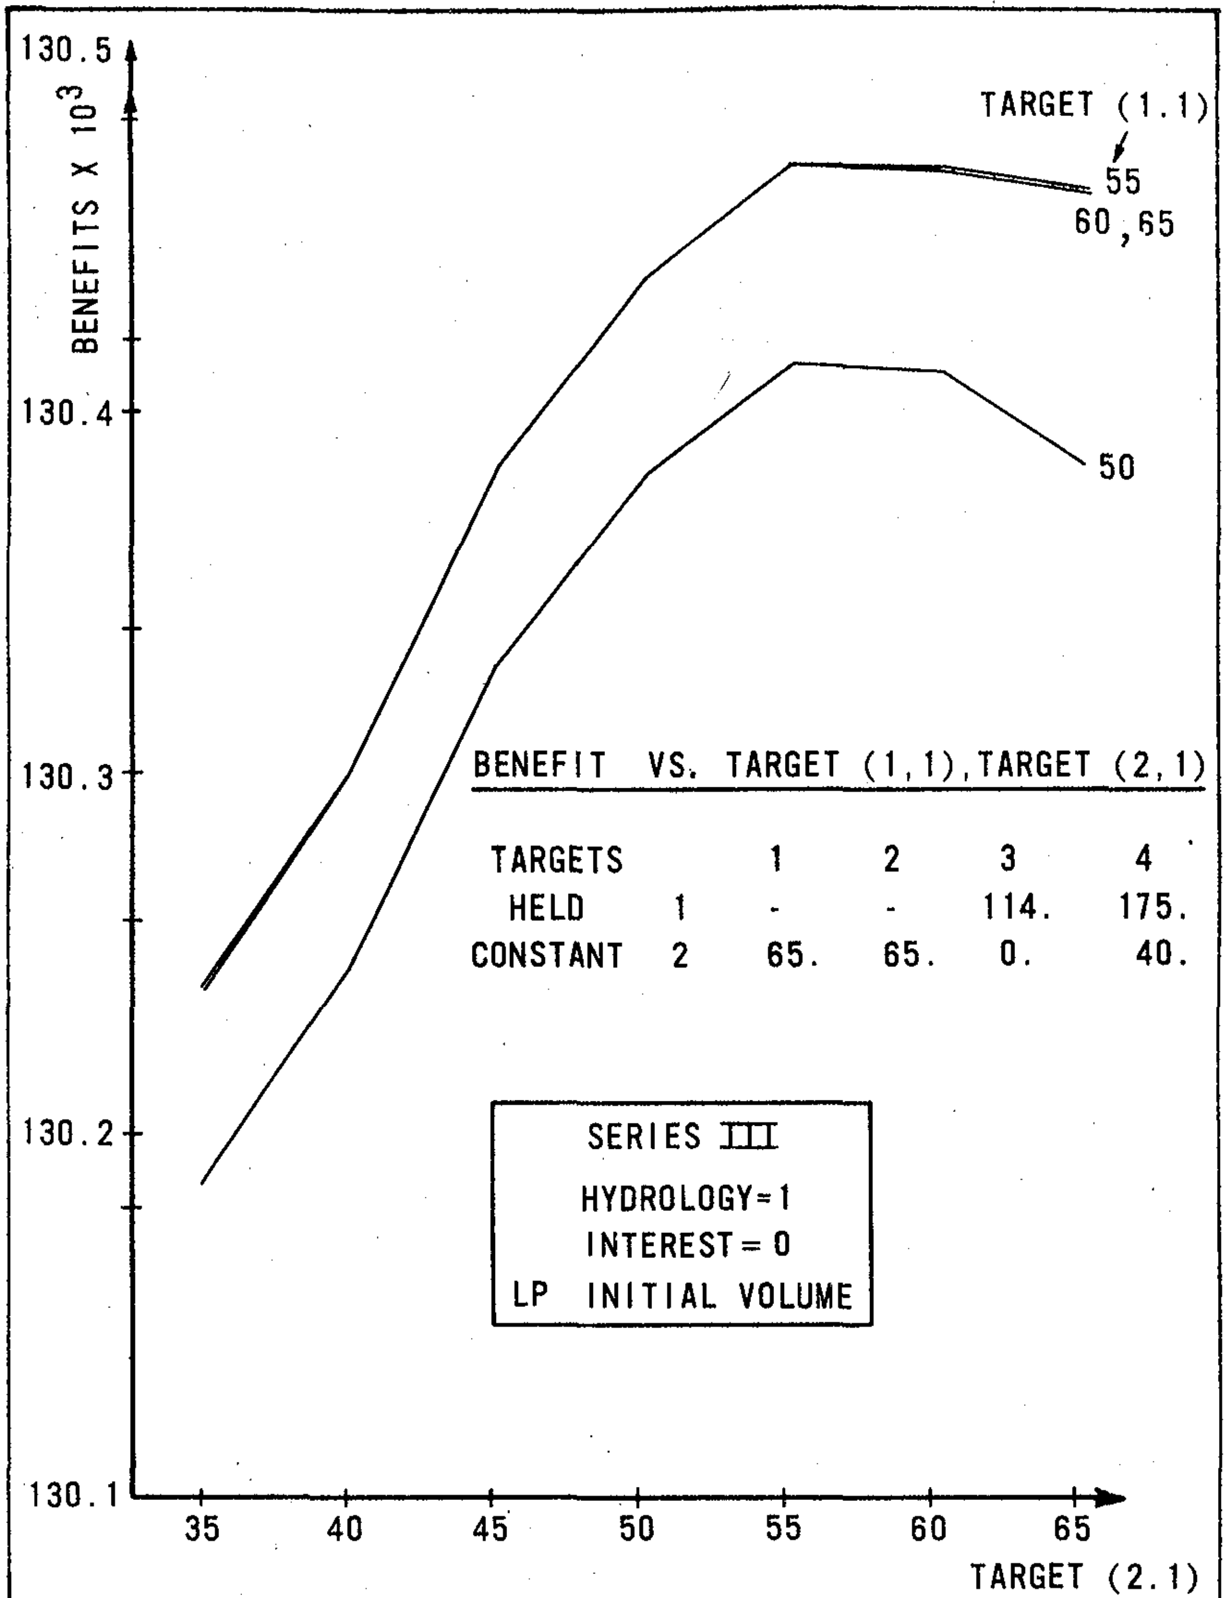

FIGURE 3.11

MARGINAL ANALSIS

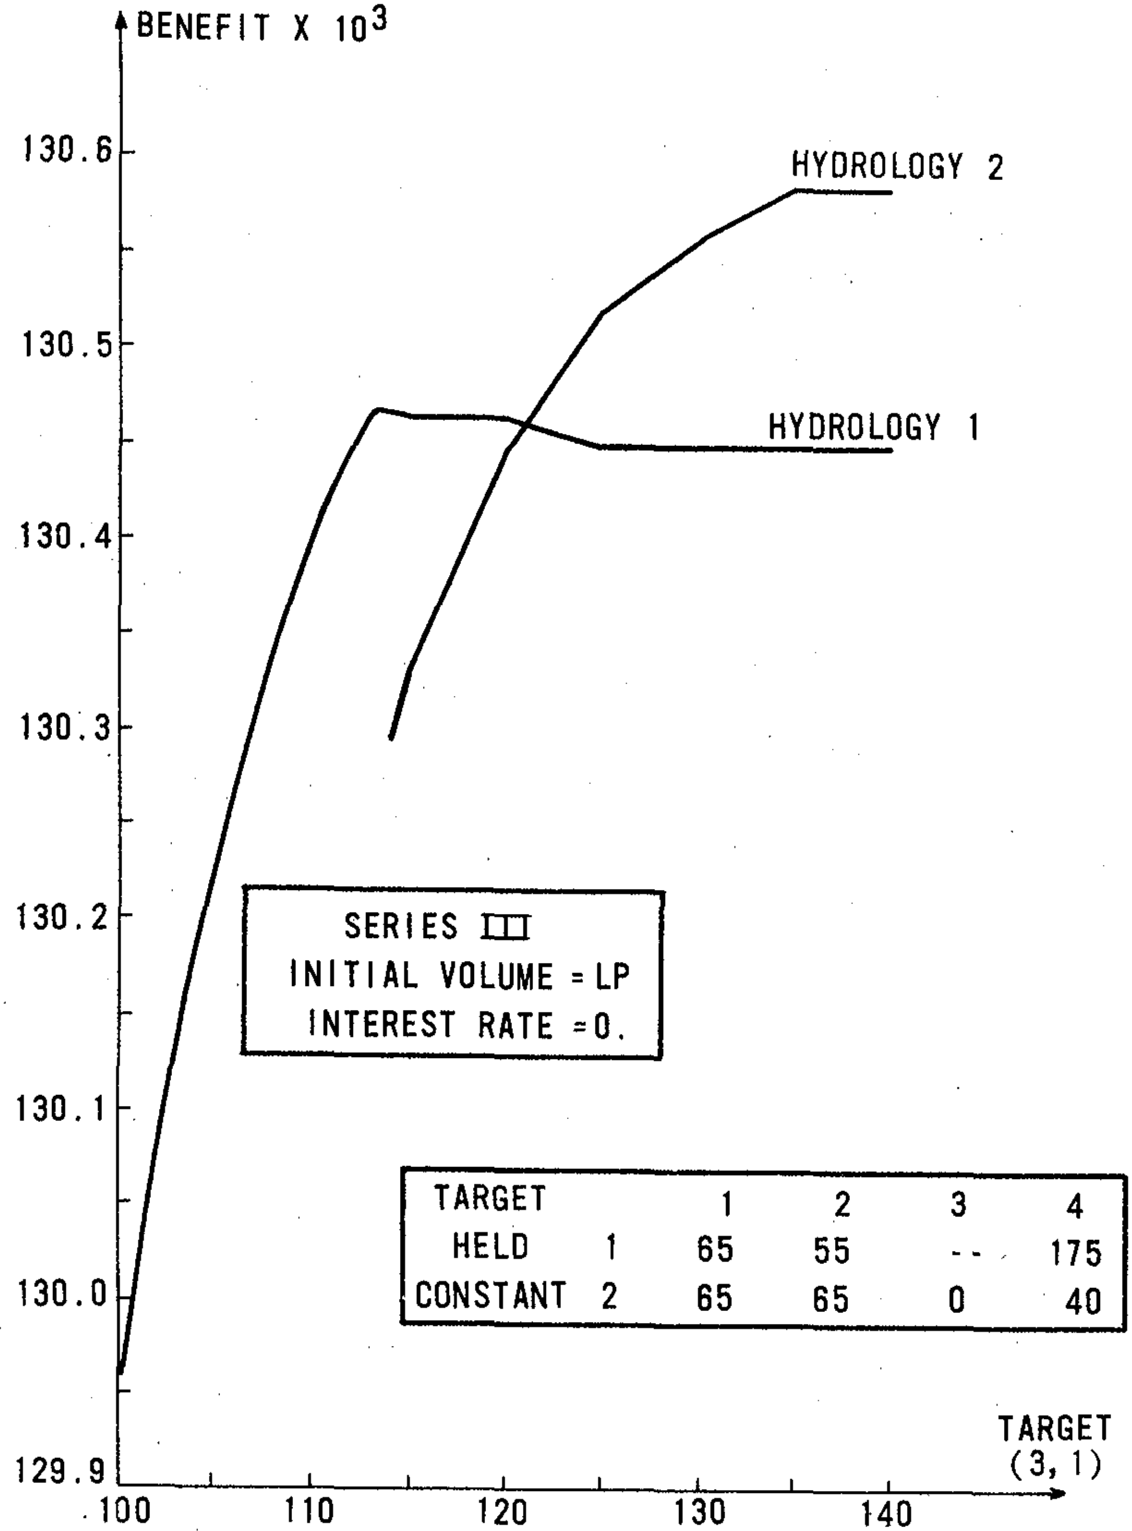

FIGURE 3.12

<u>SENSITIVITY TO TARGET</u> (3 ,1)
AT OPTIMAL POLICY

First, it serves as a check to see if the value of target (3,1) arrived at by sensitivity analysis as shown in Table 3.6c has changed by virtue of the change in targets (1,1) and (2,1) made as a result of marginal analysis. Secondly, it provides information for choice of a unique near-optimal parameter set.

The selection of such a parameter set may be cast as a problem of game theory. Synthetically generated hydrologies should be equally-likely samples from the total streamflow population. Accordingly, selection of a value of target (3,1) may be based on the principal of maximizing expected value of benefits for all hydrologies. For the results of Figure 3.12, this approach would select a level of 135 or 140 units, for an expected benefit of 130,515 units. A decision to maximize the minimum benefit obtained from each hydrology yields a target value of 121 units, at the point where the two curves cross. The associated benefit is 130,460 units. Minimizing maximum risk selects a target of 135 or 140 units. Thus, the choice will depend upon economic goals inherent in each decision criterion.

Figure 3.13 presents results of an axial analysis study. Greatest sensitivity along the axes is shown to target (3,1), a result borne out by the single-factor analysis and sensitivity analysis. Figure 3.13 shows that a significant increase in benefits is obtained by the introduction of regulation at reservoir 3 alone, and somewhat less so for reservoir 4 alone. These reservoirs serve to re-regulate upstream flow for irrigation purposes. Introduction of regulation at reservoirs 1 and 2, which regulate for power production, do not show significant differences from run of the river operation.

Axial analysis may be used to generate an initial parameter set for the LPSS process in the absence of a mathematical programming model.

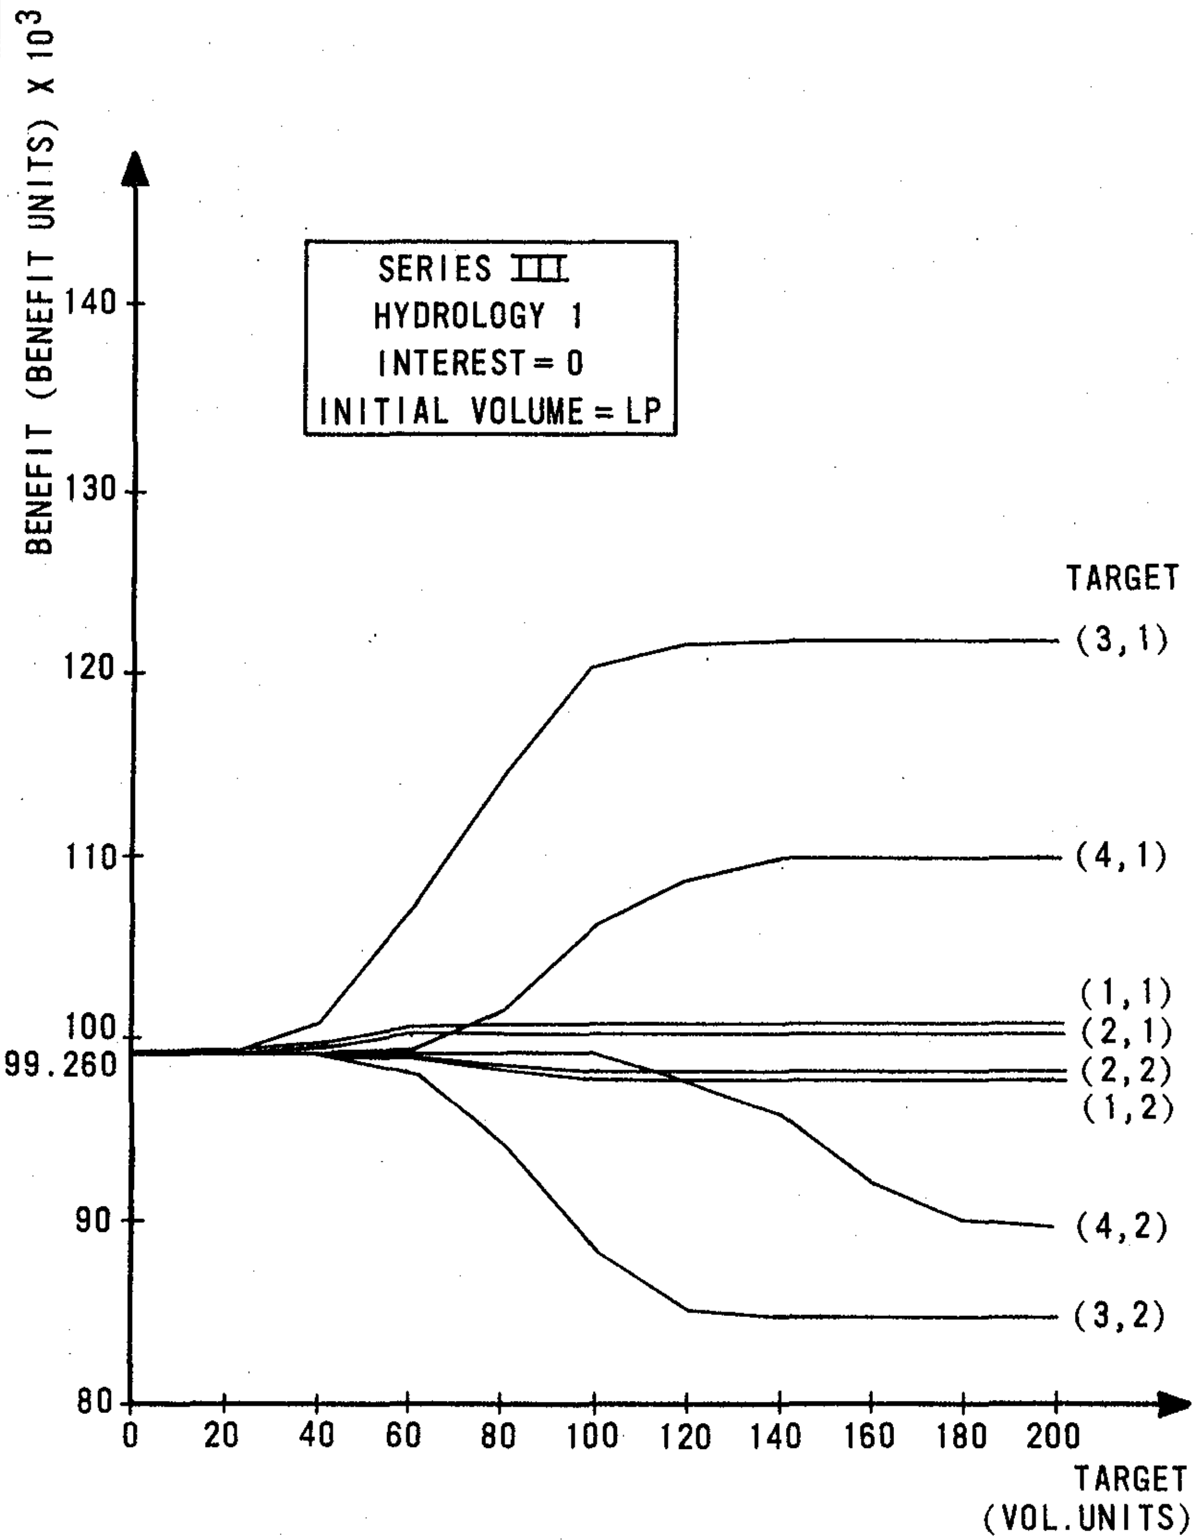

FIGURE 3.13

TARGET VALUE VS. BENEFITS (AXIAL ANALYSIS)

The value of the target that yields maximum benefit along each axis is obtained, and the values for the eight axes combined to form a single parameter set. Results of a simulation using a parameter set so-obtained yield benefits of 128,320 units for hydrology 1, which may be compared to benefits obtained for the 7 initial policies. The axial analysis parameter set is superior to the three randomly generated parameter sets, and somewhat inferior to the linear programming parameter set.

3.5.4 Series IV

Results of Series IV studies of sensitivity of near-optimal operation to capacity are presented in Figures 3.14 and 3.15, for an interest rate of). and hydrology 1. The near-optimal parameter set for hydrology 1 obtained in Series III for LP capacities is used in simulation runs with varying reservoir capacity parameter. The value of the asymptotic benefit is obtained for reservoir capacity parameters of 40 and above. A value of 1.0 for the reservoir capacity parameter corresponds to the linear programming capacities. Examination of Figure 3.14 shows that increase in benefits is close to linear from zero to linear programming capacities. Increasing capacity beyond linear programming values yields rapidly diminishing returns to scale. Since the simulation corresponds to the linear programming model in everything except stochastic effects, the diminishing returns to scale are attributed to the capability of the large reservoirs to regulate and store the occassional extremely large flows. The shape of Figure 3.14 for a non-stochastic case, where the precise inflows used in the LP model are also used each year in the simulation model, would show a completely horizontal segment past the linear programming capacities.

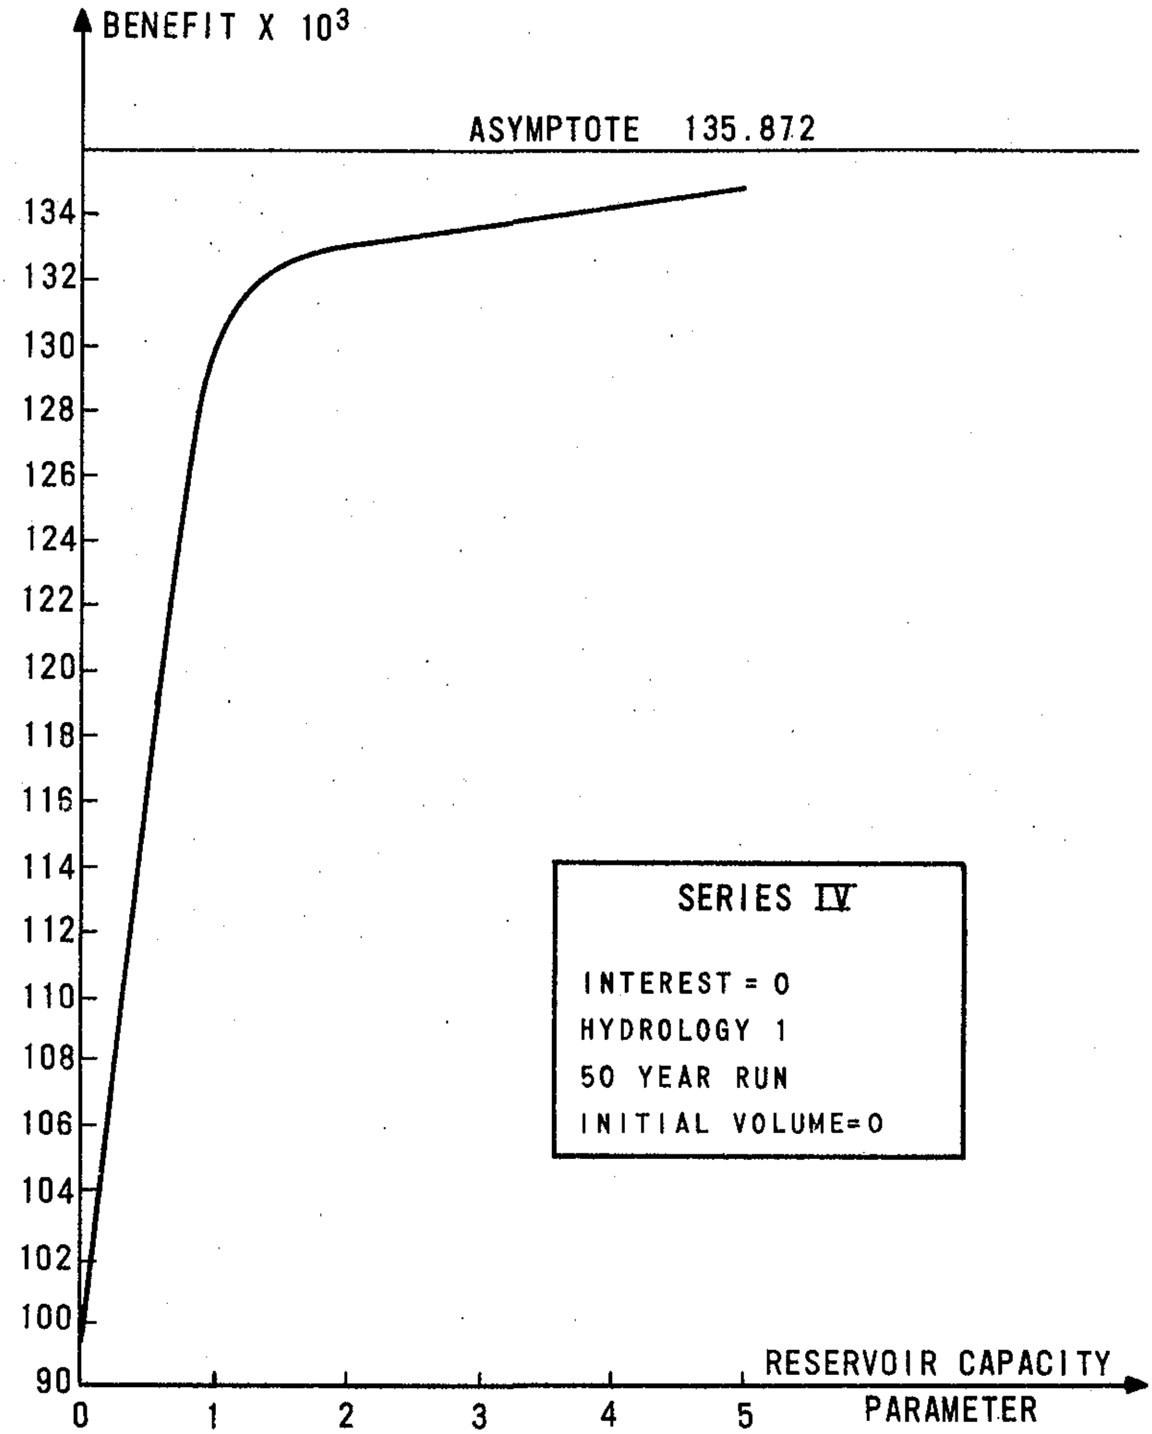

FIGURE 3.14

BENEFIT VS. CAPACITY

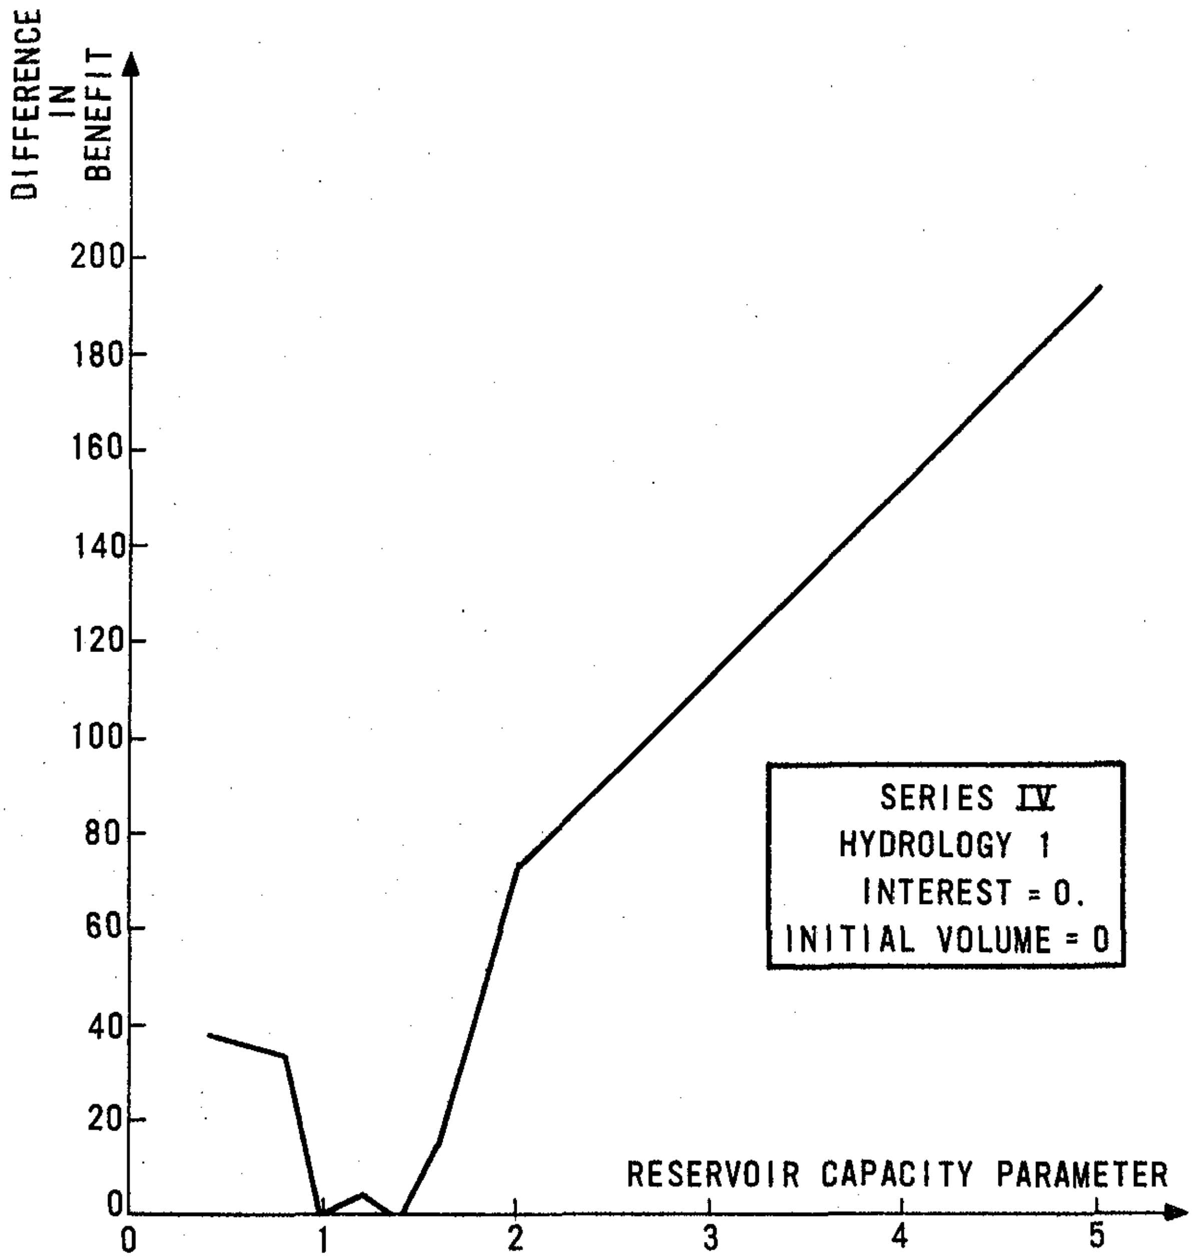

FIGURE 3.15

SENSITIVITY OF NEAR OPTIMAL OPERATION TO RESERVOIR CAPACITY

To obtain the results shown in Figure 3.15, the near-optimal parameter set was obtained for each reservoir capacity parameter by an LPSS process. For each reservoir capacity parameter value, the difference in benefits between the near-optimal parameter set and the value presented in Figure 3.14 was obtained. This difference in benefits represents the value of performing an LPSS process for each reservoir capacity parameter, rather than using the results of the single LPSS for reservoir capacity parameter of 1.0 for all other values.

Results show that differences in benefits increase as reservoir capacity parameter both increases and decreases from a value of 1.0. The return to zero difference at 1.4 units may be explained by the grid spacing selected in performing the LPSS, which may have failed to generate a slightly higher benefit by missing the best near-optimal parameter set. For most of the studies presented here, the grid spacing was taken as 5 units. A finer grid spacing may have been needed for this particular case.

Although the near-optimal parameter set does change as a function of reservoir capacity parameter, the magnitude of the change in benefits is not large. The percentage increase in benefits obtained by determining the near-optimal policy at a parameter value of 5.0 is less than .2%. This indicates that moderate changes in reservoir capacity may be made without redefining the near-optimal parameter set for each change.

3.5.5 Series V

Results for a Series V study of effect of initial volume on the near-optimal parameter set are shown in Figure 3.16. The initial volume parameter represents the percentage of total capacity present in

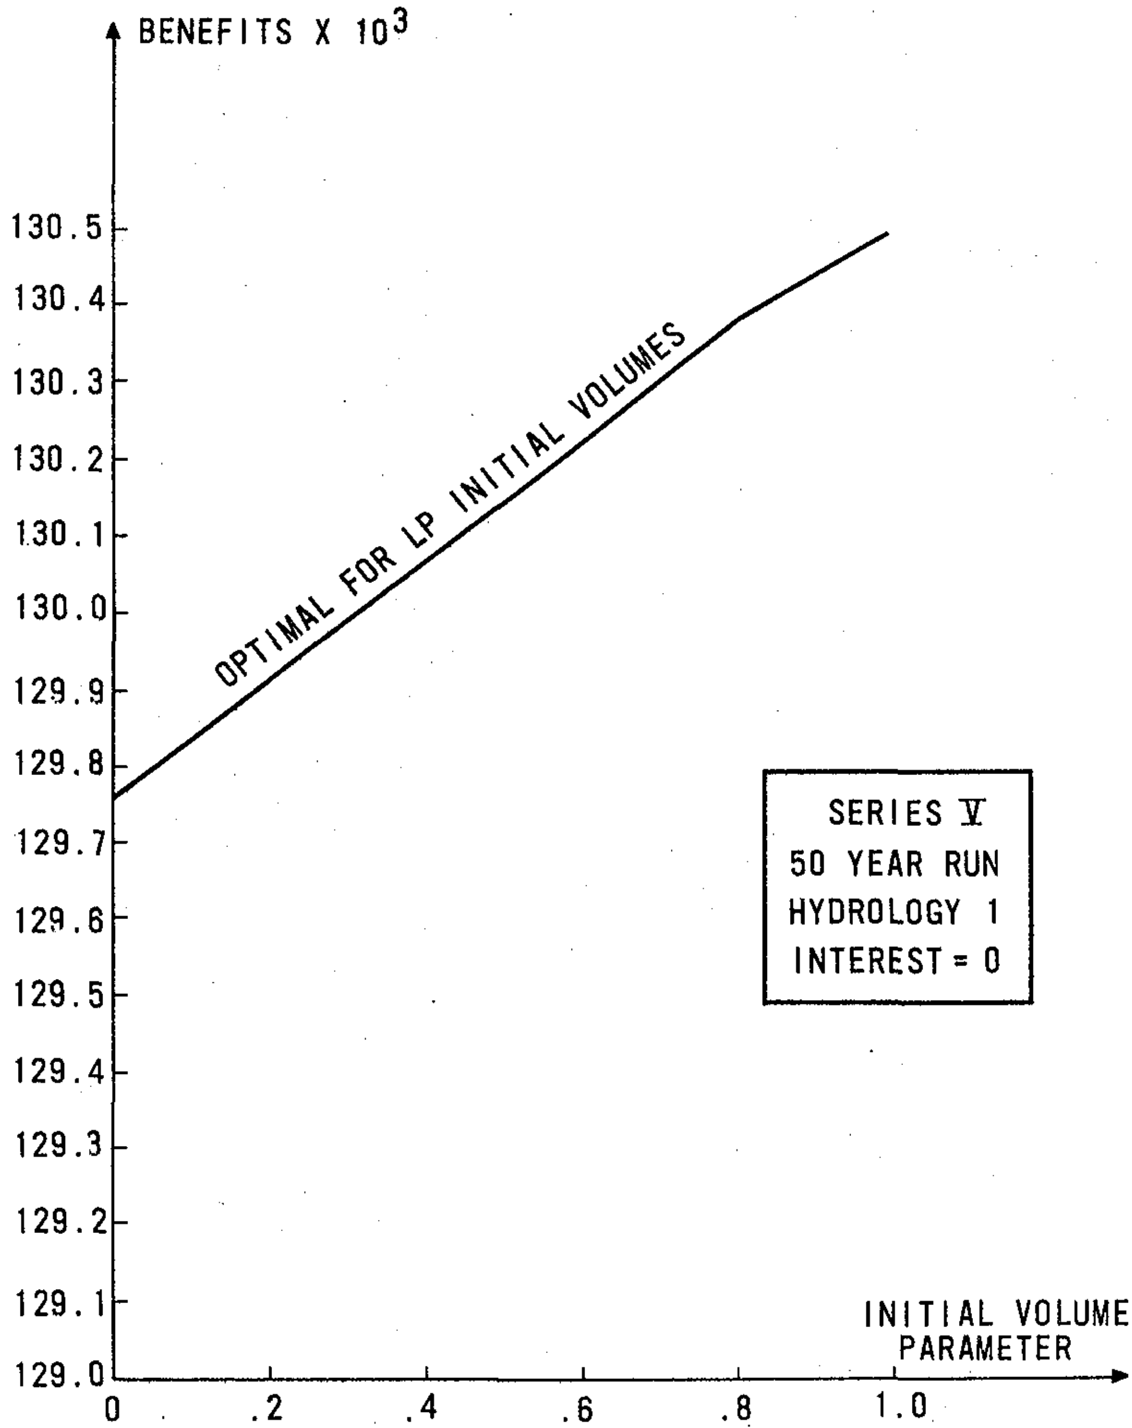

FIGURE 3.16

SENSITIVITY OF BENEFITS TO INITIAL VOLUME the reservoirs at the start of a simulation; a value of 1.0 starts the simulation with full reservoirs, a value of 0.0 starts the simulation with empty reservoirs. The near-optimal parameter set used is that developed by LPSS using LP initial volumes. Figure 3.16 presents benefits obtained using this parameter set for varying value of initial value parameter. Benefits increase linearly to a value of the initial volume parameter of .8, and at a slightly lower slope beyond that.

An LPSS process carried out for each value of initial volume parameter did not improve the near-optimal parameter set used, and no difference in benefits was found. Thus, for the 50-year duration of simulation, near-optimal parameter sets are found to be insensitive to initial reservoir volume. Similar studies were also carried out for other values of interest rate, but are not presented here, since in all cases the shape of the curve was identical to that found in Figure 3.16, and LPSS could not improve the near-optimal parameter set.

The studies of Series V, performed for a 50-year simulation, are not adequate to determine the short-range influence of initial volume. Later studies (Series VII) indicate that initial volume effects persist for 2 to 5 years depending on reservoir capacity, indicating that sensitivity of operation to initial volume would more likely be encountered with simulation runs of short duration. The 50-year duration tends to mask whatever effect initial volume may have.

3.4.6 Series VI

No results are presented for Series VI studies, since no sensitivity of the near-optimal parameter set to interest rate was encountered, in spite of the large effect of interest rate on benefits, as shown in Figure 3.7. Results of dynamic programming studies of reservoir operation (13) clearly indicate that a significant sensitivity of operation to interest rate does exist. The present model appears inadequate to point this out, principally by virtue of holding the policy parameter set constant throughout the 50-year simulation run. A more flexible operating policy, with time-varying target releases, might be expected to indicate some sensitivity to interest rate. Similarly, a 12-period model rather than the 2-period model used would be expected to show a greater degree of interest rate sensitivity. Non-linear benefit functions rather than the linear structure used may also cause response to interest rate.

3.5.7 Series VII

Results of SEries VII studies of the persistence effect of initial volume differences are presented in Figure 3.17, as a plot of difference in benefits obtained over the duration of simulation between initially full and initially empty reservoirs, against duration of simulation. Reservoir capacity is taken as a parameter, with LP capacities and ten times the LP capacities being investigated.

Simulation runs of varying duration were performed for the two cases of initially full and initially empty reservoirs. Benefits for the two cases were subtracted for each duration of simulation, and plotted in Figure 3.17. The final horizontal portion of each curve indicates a constant difference in benefits between the runs for initally full and initially empty reservoirs. This constant difference in benefits indicates that the effect of differences in initial volume on the operation of the system has been damped out. Thus, for the smaller reservoir capacities, persistence effects last for 5 periods, while the larger reservoir capacity shows persistence effects for 12 periods. It must be noted that the larger reservoir has a correspondingly greater amount

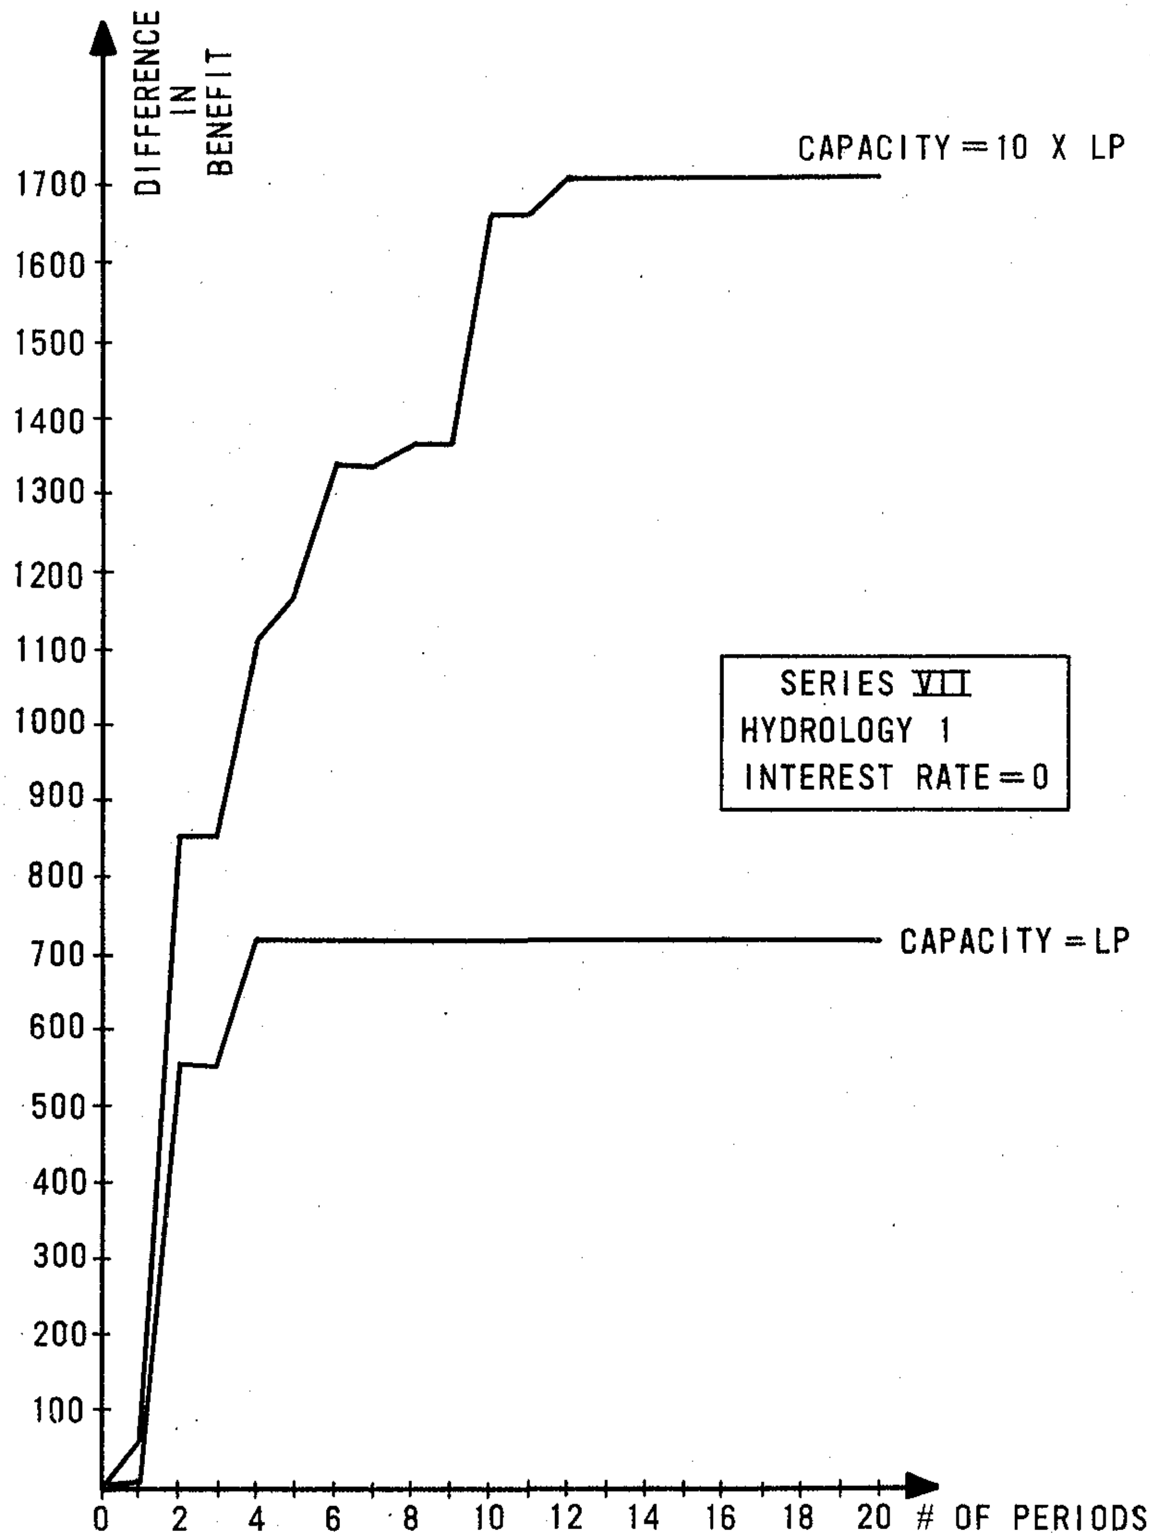

FIGURE 3.17

PERSISTENCE OF INITIAL VOLUME EFFECTS

of water available when initially full than the smaller reservoir.

Intermediate horizontal portions of the curve are due to the fact that irrigation benefits are calculated for every 2-period year, rather than for every period as power benefits are. Consequently, water supplied beyond the power capacity to satisfy irrigation requirements is reflected in benefits only every two periods.

Results of Series VII indicate the appropriate planning horizon that can be used in developing a state-sensitive policy. Since initial state effects disappear in terms of operation after the persistence period, simulation runs with durations equal to or less than the persistence period are expected to show some sensitivity of operating parameters to initial volume. As will be recalled, this was not the case for the 50-year simulation runs of Series V.

With the shorter simulation runs, near-optimal parameter sets can be obtained for different initial volumes. Consequently, a table of near-optimal parameter sets vs. state can be built up. Such a table could be used to yield a state-oriented rule for simulation by selecting in any period the appropriate near-optimal parameter set for the current system state. The feasibility of such a process depends upon the degree of sensitivity of operating parameters to state. High sensitivity would require an elaborate search process to establish such a state-oriented policy.

3.5.8 Series VIII

Two sources of discrepancy are identifiable between corresponding simulation and linear programming models. The introduction of non-linear functions in a simulation model, where only linear functions are permissible in the corresponding linear programming model, is one source. The second discrepancy comes about from the difference between the stochastic effects present in the simulation model and the fixed hydrology of most simple linear programming models.

For the present case, the simulation has been constructed with linear functions only. The only difference between the simulation and linear programming models is due to stochasticity. This difference produces the difference in operation noted above, so that the linear programming targets are not directly applicable as the near-optimal targets for the simulation. To eléminate differences in operation and compare the two models for the effects of stochasticity alone, each model may be run with zero reservoir capacities. Run of the river benefits can be compared to determine the faithfulness of the models. For zero reservoir capacities, annual benefits obtained from the fixed hydrology of the linear programming model are equal to 2422 units per year. Run of the river operation for a 50-year simulation run and 0.0 interest rate using hydrology 1 yields a total benefit of 99211 units, or 1985 benefit units per year. The difference between the two figures may be considered as the cost of stochasticity. Capacity constraints in the simulation model limit the benefits that can be obtained from excess flows, and no regulation capacity exists to ameliorate low flow conditions.

For the case of simulation models with non-linear functions, the effect of stochasticity can be eliminated in the simulation by artificially constructing a 50-year hydrology such that in each year, the inflows are equal to the mean values used in the linear programming model. The hydrology would then be a 50-year recurrence of the mean inflows, and would correspond precisely to the situation modeled in the linear programming case, provided that interest rate is held at zero. Benefits obtained from the simulation with such a hydrology would serve to isolate the effects of non-linearity while eliminating effects of stochasticity.

3.6 Conclusions

Figure 3.18 represents schematically the benefits derived from the various stages of the LPSS process. The benefits due to the best random starting point, the parameter set resulting from axial analysis, and the linear programming parameter set are indicated along the benefit axis. Benefits associated with parameter set 7 are shown slightly to the right of the benefit axis, to indicate that policy 7 is not one of the "pure" initial policies. Rather, the parameter set is derived after a degree of search has been carried out.

The axial analysis is seen to be a significant improvement over the best of the three random parameter sets, while the linear programming presents an improvement over the axial analysis. The scale of Figure 3.18 is somewhat misleading, since the benefit axis does not start from 0. benefit units. Consequently, the percentage change in benefits is not large, with an increase of 7% for LP over random parameter sets.

Although the random starting point does not in itself provide high initial benefits, it allows the generation of parameter set 7, which is seen to be an improvement over the linear programming set. Starting from the linear programming set, incremental benefit due to parameter set 7 is 339 benefit units, an increase of .26%. The addition of a single-factor analysis augments benefits by 39 units, while the next step of sensitivity analysis to target (3,1) increases benefits by 6 units. The final stage of marginal analysis further increases benefits by 5 units. Thus, the principal incremental benefit

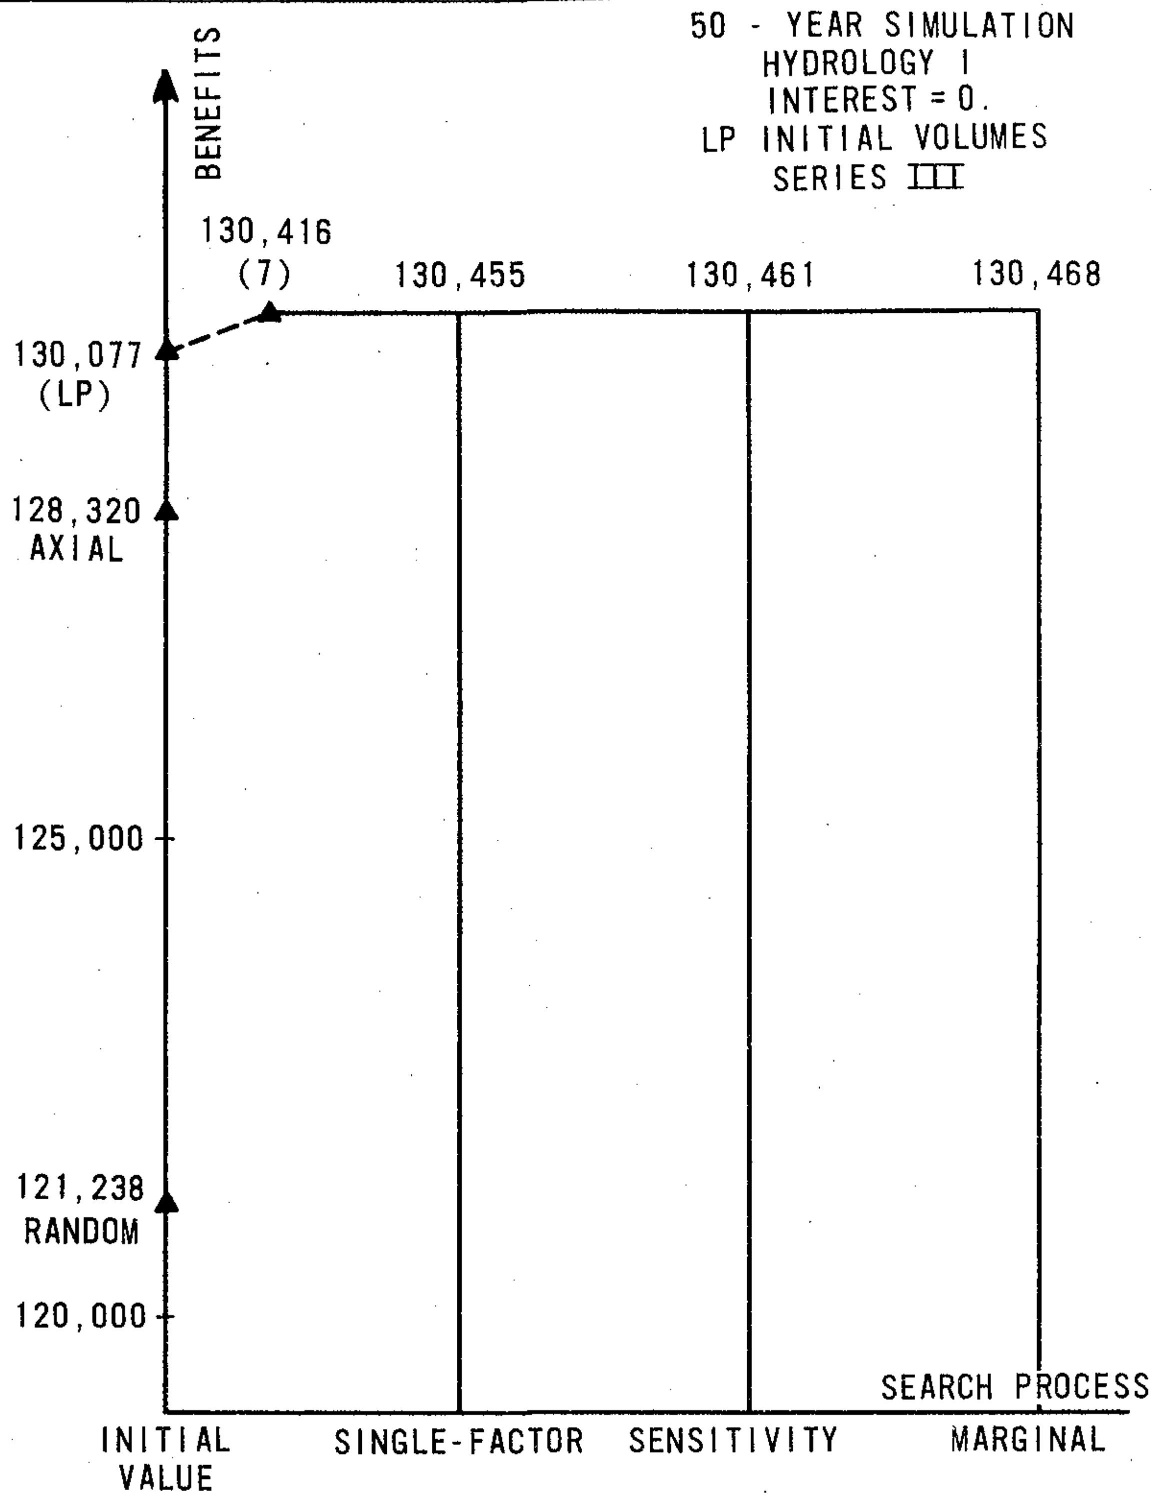

FIG 3.18

INCREASE IN BENEFITS FROM LPSS PROCESS

is obtained through application of the linear programming technique.

The above results show that the incremental value of each successive search technique decreases rapidly. The total LPSS process therefore converges rapidly from an appropriate initial value. Consequently, the LPSS process is seen as a practical method of obtaining parameters for operating rules. Use of linear programming models to provide initial values is seen to be the most important part of the search technique in terms of incrementing benefits. Thus, it is an integral part of the search process, and linear programming models should not be considered as alternative search techniques, rather as a fundamental part of a total search process.

Of the search techniques used, single factor analysis is seen to most rapidly indicate appropriate directions of change. The combination of the three techniques of single factor, sensitivity, and marginal analysis was found totally adequate to determine the near-optimal parameter set.

It is recognized that the advantages of the LPSS process shown here in part stem from the fact that the artificial model was constructed to conform closely to the linear programming model, and as a consequence, linear programming results yield good starting values for the search process. Further, the relatively simple nature of the system allowed rapid evaluation of a majority of the operating policy parameters. This will not always be the case, and the LPSS process may take considerably longer to converge than the process shown in Figure 3.18. The technique of LPSS is sufficiently simple and flexible however to be adaptable to longer search processes with no conceptual change, and as such is seen to be useful for more than the simple model presented here.

Use of the LPSS process should be controlled by a plot of the type of Figure 3.18. By means of such a plot, trade-offs between further search and increase in benefits can be rapidly evaluated, and the decision to continue or terminate a search process be made with reference to incremental benefits of each search technique.

CHAPTER IV

THE MAULE BASIN MODEL

4.1 Introduction

The Maule basin of central Chile was selected in 1963 to be used as a test case for the research on systems methodology. Various studies have been carried out based on data from the Maule region. Leonvendagar and McLaughlin (16) have developed a simulation model of the basin, and Ibanez developed computer programs for the operation of the basin according to four different operating policies (15). Wallace used data from the Maule basin for his linear programming studies (5), and Poblete has produced a detailed study of linear programming models of the Maule basin (20). The existence of the simulation model with external operating rules and the existence of mathematical programming models for the basin dictated the choice of the Maule basin to investigate the feasibility of the LPSS technique in the context of an existing basin.

The large-scale long-term nature of the investigation of the Maule basin makes the staging in time of the information received an important factor in this study. During the course of all the studies of the Maule, both new basic data and further sophistications of the mathematical models were being introduced. Each study would start with the data available at the time, and attempt to update in so far as possible, as time progressed. This staging of information is typical of "real-world" engineering situations, and the course of such staging and its pertinence to the present investigation will be included as a basic part of the research.

4.2 Description of the Basin

Detailed descriptions of the Maule basin are found in references (16) and (20). A brief outline will be given here, taken mainly from Ibanez.

Figure 4.1 shows the location of the Maule basin in Chile, and Figure 4.2 is a map of the basin itself. The total area of the basin is approximately 21,500 km^2, and the population as of 1965 was 530,000. Talca is the principal city of the basin with a population of 80,300. The principal river, the Maule, has two main tributaries from the north, the Cipreses and the Claro, and two from the south, the Melado and the Loncomilla. The primary economic activity of the area is agriculture, and the primary use of water is for irrigation and the generation of hydroelectric power, which is exported to the national power grid. To simplify the model somewhat, only the upper basin, containing the major features, was modelled.

The principal elements and the configuration of the model are shown in Figure 4.3, and identified in Table 4.1. Two major reservoirs, Invernada and Laguna Maule, are in existence, as are two power plants, Isla and Cipreses. Two reservoirs, Colbun and Guaiquivillo, are proposed, as are two power plants, Central Maule and Central Colbun. The existing and proposed capacities of these elements are given in Table 4.2.

4.3 The Simulation Model

The simulation model used in this study is essentially that developed by Leonvendagar and McLaughlin (16) with certain changes in coding, input-output, data storage, and treatment of the operating rules. Coding is in Fortran IV, and the model was implemented on the IBM 360/65 computer.

The model uses a monthly period, with an external operating rule. A simulated hydrology of the required length is generated for all rivers

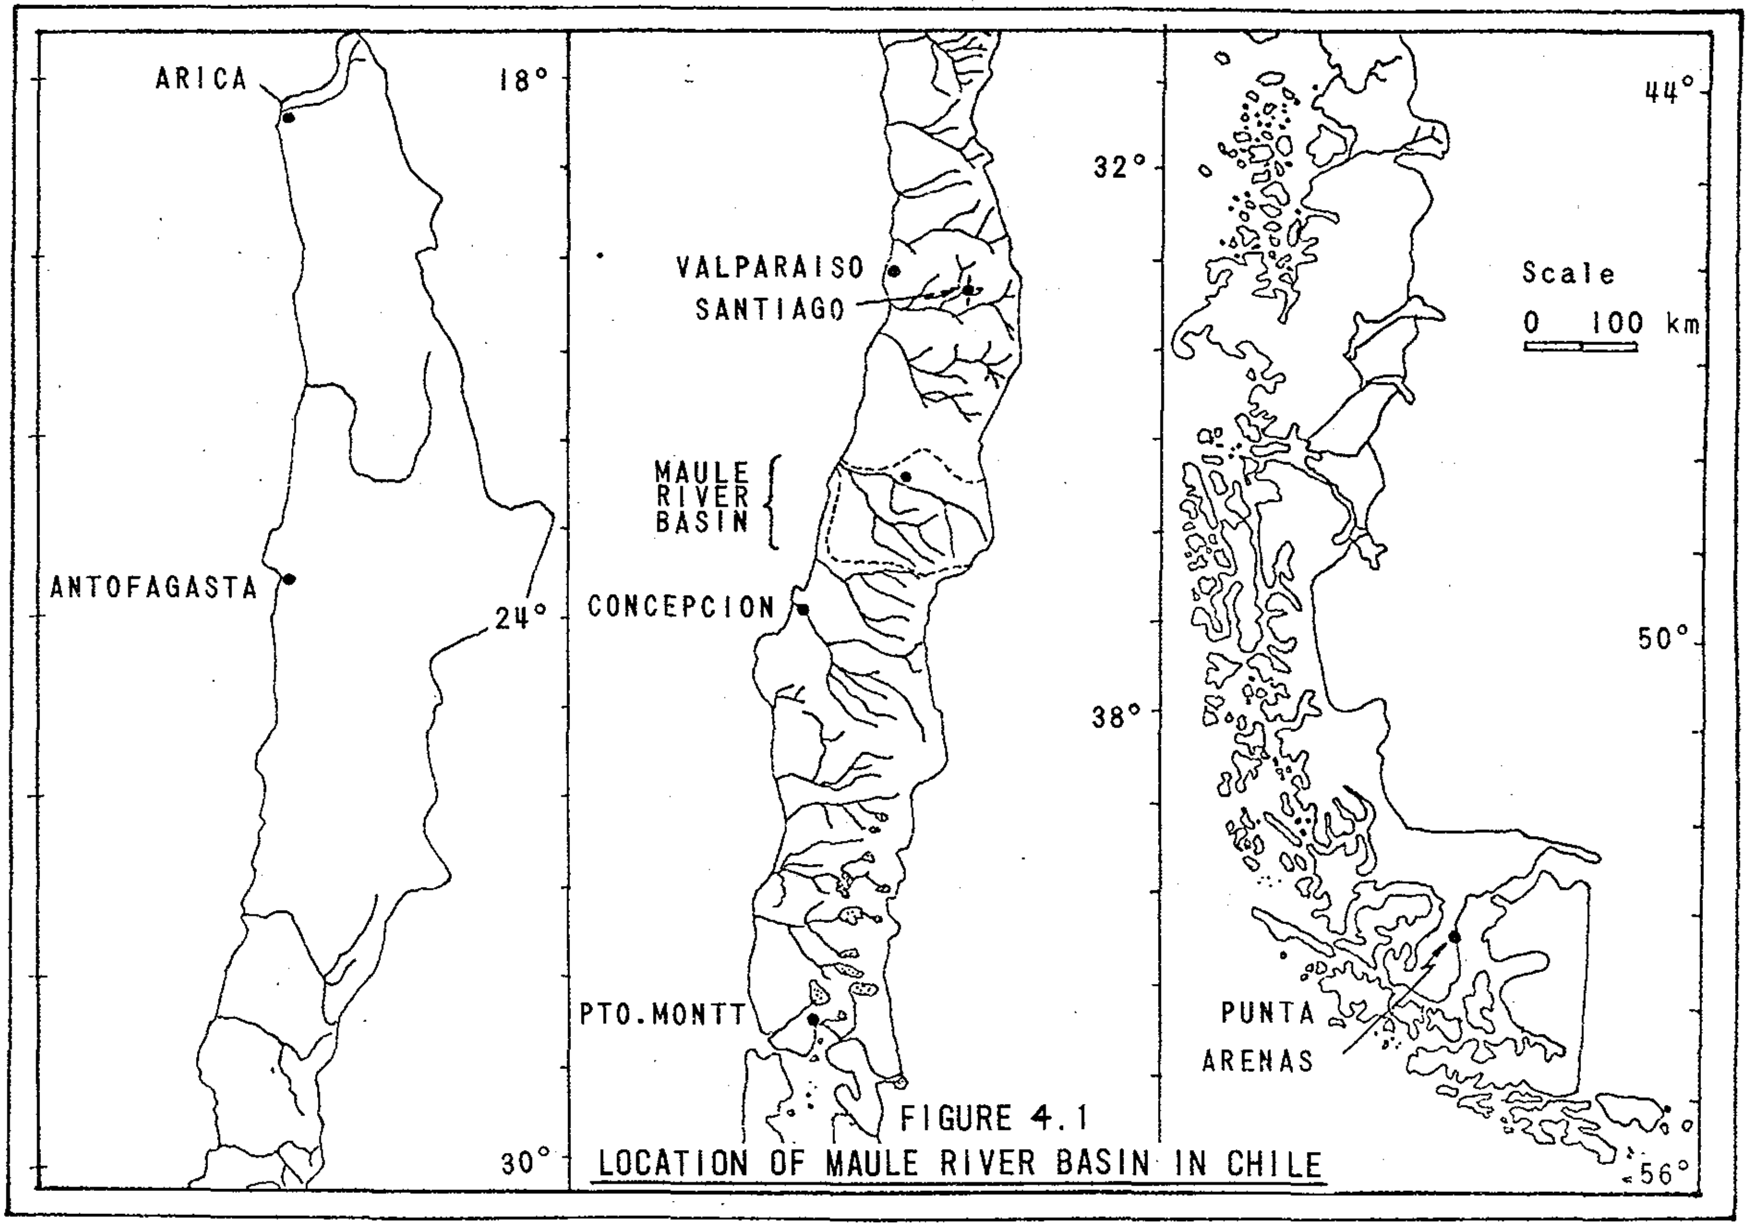

FIGURE 4.1
LOCATION OF MAULE RIVER BASIN IN CHILE

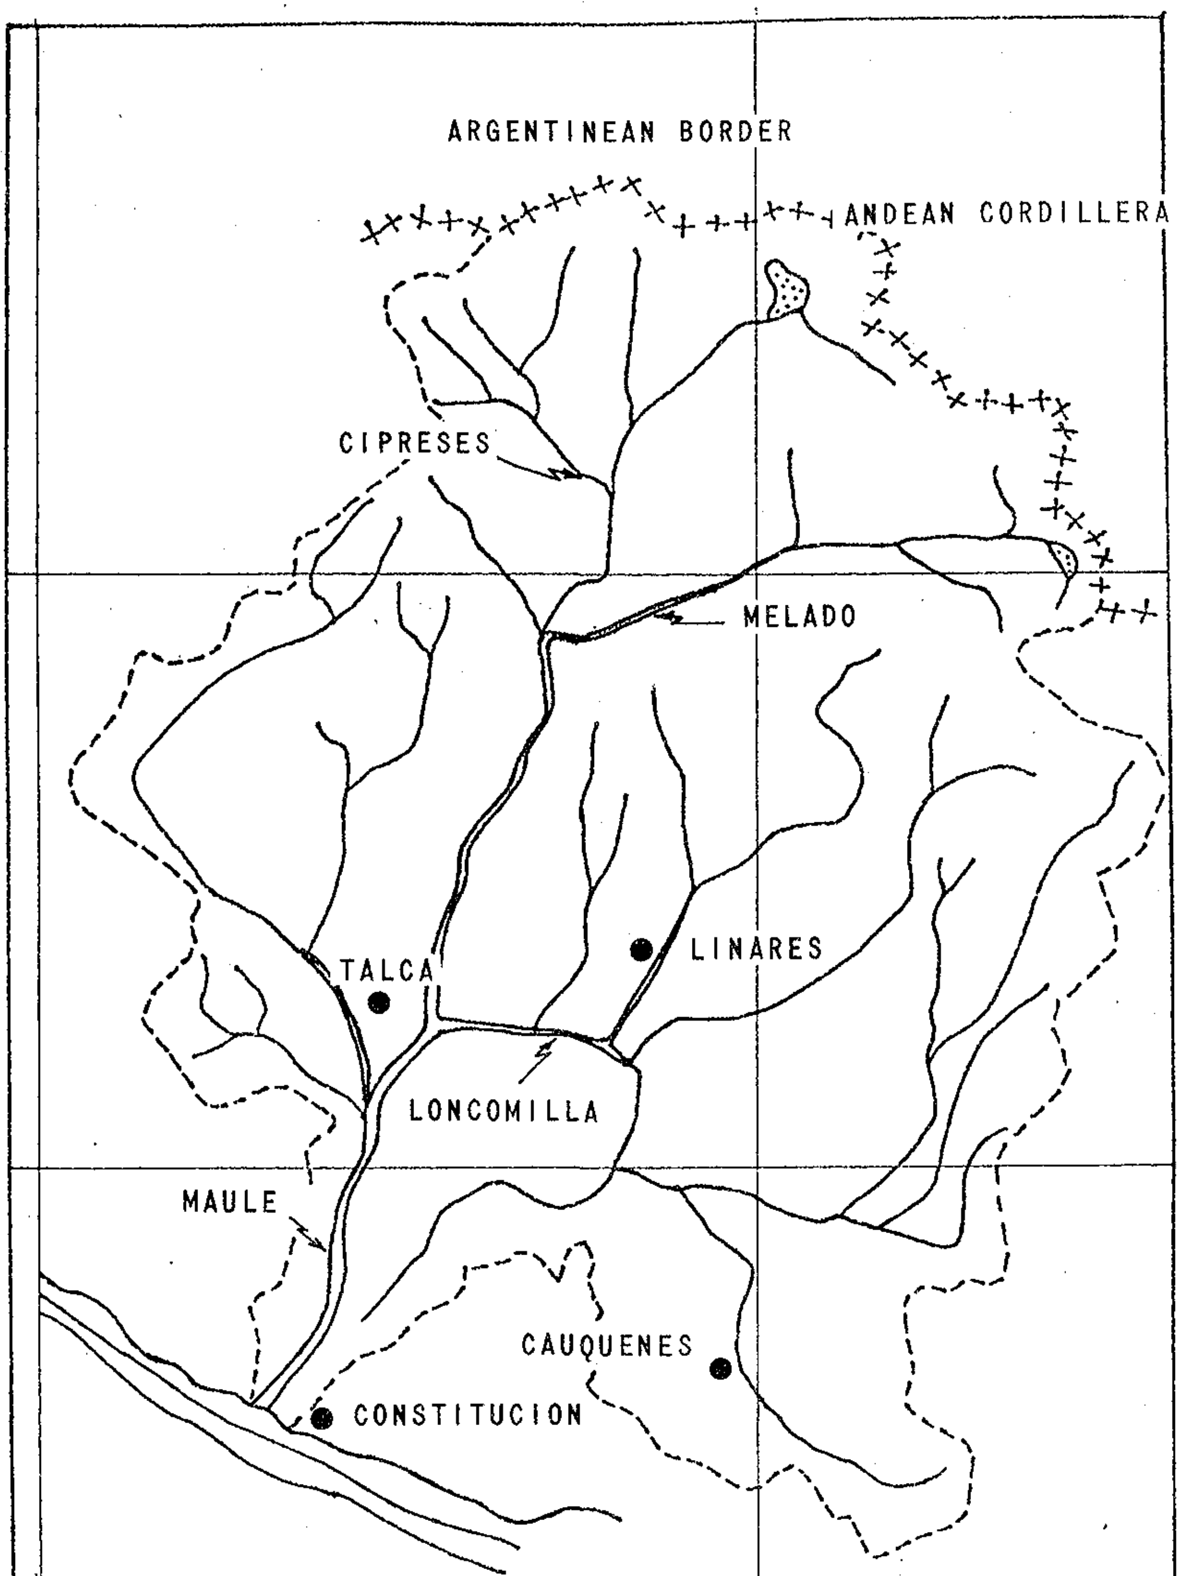

FIGURE 4.2
MAULE RIVER BASIN

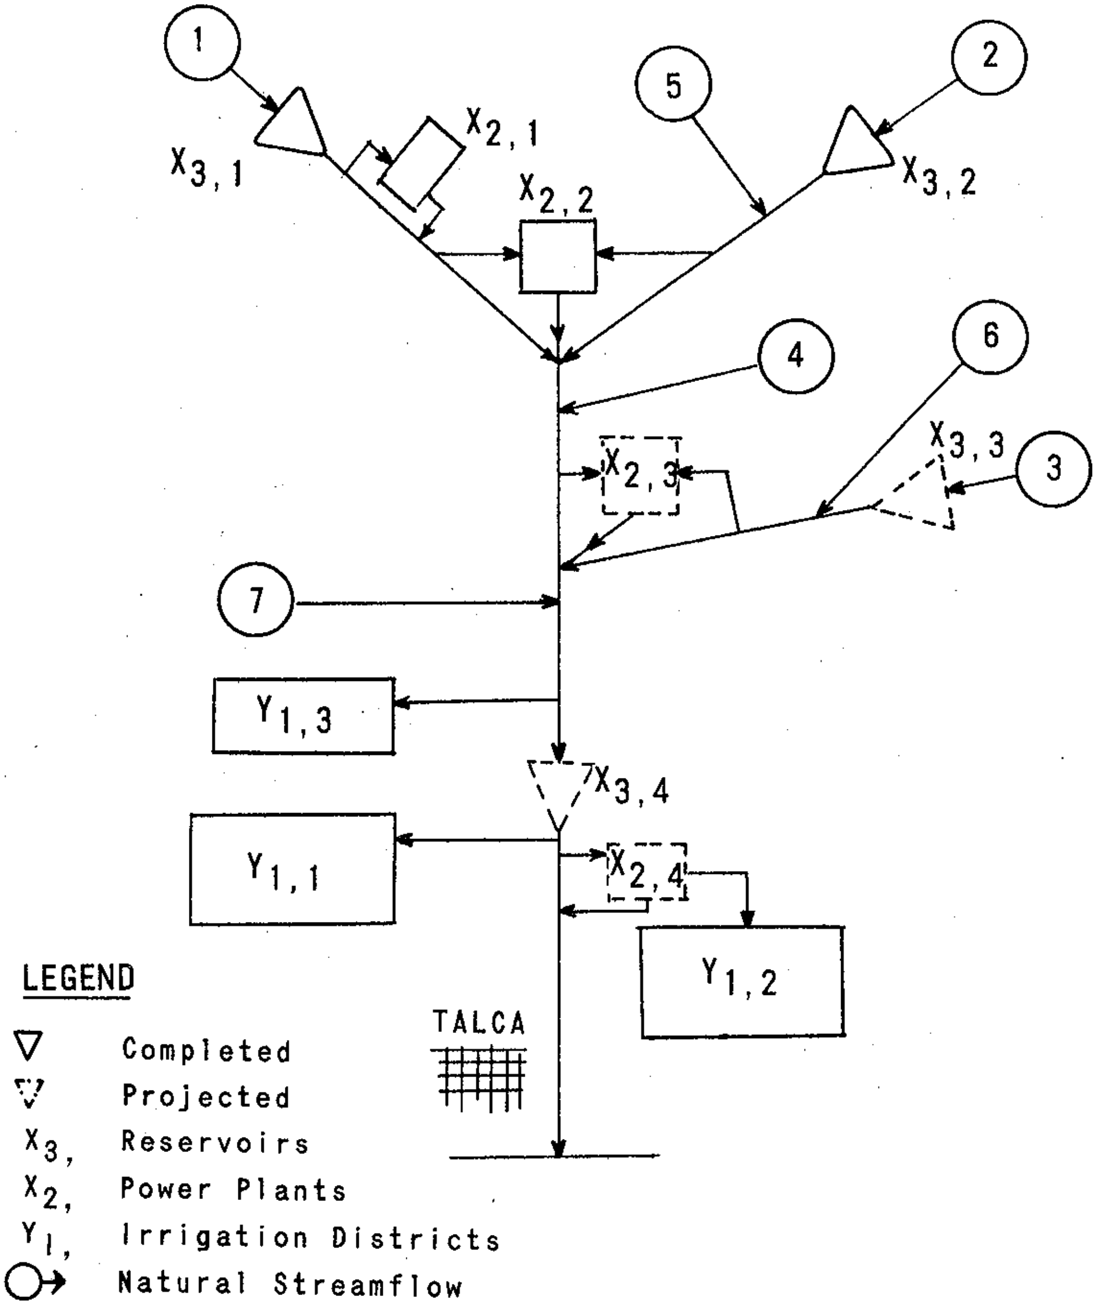

FIGURE 4.3

SCHEMATIC REPRESENTATION OF REGION USED IN THE MODEL

TABLE 4.1

ELEMENTS OF THE MAULE MODEL

$X_{3,1}$	= Reservoir Laguna Invernada (completed)
$X_{3,2}$	= Reservoir Laguna Maule (completed)
$X_{3,3}$	= Reservoir Guaiquivillo (projected)
$X_{3,4}$	= Reservoir Colbun (projected)
$X_{2,1}$	= Power plant at Cipreses (completed)
$X_{2,2}$	= Power plant at Isla (completed)
$X_{2,3}$	= Power plant at Maule (projected)
$X_{2,4}$	= Power plant at Colbun (projected)
$Y_{1,1}$	= Irrigation area Maule Norte Bajo
$Y_{1,2}$	= Irrigation area Maule Sur
$Y_{1,3}$	= Irrigation area Maule Norte Alto
1	= inflow to Laguna Invernada
2	= inflow to Laguna Maule
3	= flow in Melado River
4	= flow in Colorado River
5	= flows in Rivers Campanario and Puelche
6	= flow in De la Puente River
7	= flow in Claro River
$Y_{2,1}$	= Power from Cipreses plant
$Y_{2,2}$	= Power from Central Isla
$Y_{2,3}$	= Power from Central Maule
$Y_{2,4}$	= Power from Colbun plant

TABLE 4.2

SIZES OF ELEMENTS IN THE MAULE BASIN

Reservoirs	Capacity (10^6 cubic meters)
Maule	1430.0
Invernada	157.5
Guaiquivillo	1500.0 (maximum)
Guaiquivillo -revised	460.0 (maximum)
Colbun	2500.0 (maximum)

Power Plants	Installed Capacity (Megawatts)
Isla	68
Cipreses	101.4
Maule	480 (projected)
Colbun	560 (projected)

at the start of each simulation run, with provision for cross and serial correlation and variable lag. The shape of the probability distribution may be either normal or log-normal.

Features of the model include a consideration of evaporation and filtration, consideration of variations of power production with head, and a highly sophisticated treatment of irrigation, including evapo-transpiration, a variety of crops, and crop die-off due to insufficient water supply.

Benefits are calculated after each month for power benefits, and at the end of each year for irrigation benefits. Up to 30 different interest rates may be handled for each simulation run. At the end of the run, the costs are apportioned over the period of simulation and subtracted from accumulated discounted benefits, to obtain net benefit discounted to present value.

A large number of options for output allow for monthly, yearly, and run summary output of physical and economic parameters as desired by the user. Economic information is given for each interest rate.

Running times on the IBM 360/65 under OS/360 are highly dependent on the quantity of output requested. A forty-year simulation run with output consisting of run summary economic information for 8 different interest rates requires approximately 1 minute of computer time. A 40-year run obtaining monthly, yearly, and run summary information on all physical and economic options available to the user requires 14 minutes for three different interest rates.

4.4 Operating Policies

4.4.1 Introduction

The operating policies used in this study are the four policies developed by Ibanez, two policies derived from the policies of Ibanez, and one policy developed by the author. The first four by Ibanez were modified to make them compatible with the revised simulation. Ibanez developed these four policies, but did not have opportunity to do more that preliminary testing with any of them. Thus, no information on their relative merits for a full-scale simulation was available, and these studies were included in the present investigation as Part I of investigations into the operating rules. Each operating policy will be described briefly below. Certain considerations hold true for all operating policies used. They are the following:

1) Reservoirs in the upper basin will produce as much power as they can subject to specified system targets and available capacity of penstocks and turbines, in all cases except the legal operating rule, where a portion of the water in one of the reservoirs (Maule) is reserved for irrigation by a legal agreement actually in force in the basin.
2) Colbun reservoir, in the lower basin, will be used to re-regulate the flows from the three reservoirs in the upper basin, with its main purpose the supply of irrigation water. The power plant at Colbun is subject to the availability of water for irrigation requirements.

Therefore, the significant operating decisions pertain to the upstream reservoirs, since Colbun is used for re-regulation of the available water from upstream.

4.4.2 Standard Operating Policy

The standard operating policy is similar to the one described in Chapter III. Target values for each of the 12 months for each of the three reservoirs are required as auxiliary input to the simulation model when this operating policy is used.

As developed by Ibanez, the following steps are carried out in the monthly operation of the operating rule:

1) The available water, equal to the water in storage plus the inflow in the current month, is determined for the three reservoirs with natural (unregulated) inflow.

2) Based on the available water during the month, a release is made if necessary to avoid spilling, taking into account filtration and evaporation losses. These releases are termed basic releases.

3) After the basic releases are determined, further releases are made in each reservoir according to a standard operating policy, with targets varying from month to month for each of the three upper basin reservoirs. This release is made in addition to the basic release, and in this form of the operating policy the basic release doe not contribute to satisfying the target demand.

4) The amount of water available to Colbun reservoir is determined from the total releases for the upper basin. This water is supplied to the irrigation areas, and if the amount of water is insufficient to satisfy all requirements, the water is apportioned among the irrigated areas in proportion to the requirements (a linearity assumption). Irrigation water that may be passed through the turbines after the irrigation requirements are met is used to generate power at Colbun power plant, but no releases are made from Colbun reservoir specifically

to generate power.

Thus, the standard operating policy operates the three reservoirs in the upper basin to supply the target demands if possible, after releases have been made to avoid unnecessary spills. The discretization procedure and the coding imply that water released as a basic release to avoid a spill cannot be used to satisfy the monthly target for the reservoir, but downstream benefits can accrue to these basic releases. Colbun operates to satisfy the irrigation requirements in so far as possible, and the power generated by Colbun is due only to irrigation releases and unavoidable spills, not to specific releases for power. Thus, an uncoupling of the upper and lower regions of the Maule basin is provided by the operation of the proposed Colbun reservoir and power plant.

4.4.3 Revised Standard Operating Policy

The policy described in the previous section is unsatisfactory for use with linear programming models to generate target values, since the basic releases are not used to satisfy the target demand. The above policy has been revised by the author to first calculate the basic releases as before, but these basic releases are then subtracted from the appropriate target value for the month and reservoir. The standard operating policy is then applied with the reduced target level as before. This revision accurately reflects the character of target releases generated by linear programming models, which are total releases, comprising both controlled and uncontrolled flow. The original policy of Ibanez is clearly wasteful of water if the target values reflect downstream needs. This new structure of the operating policy is termed the revised standard operating policy.

4.4.4 Legal Operating Policy

In the Maule basin, a legal agreement is in force between the power and irrigation authorities. The terms of this agreement, cited in reference (21), are:

1) The Maule reservoir will be divided into two zones, the upper of 900 million m^3, and the lower of 400 million m^3.

2) When the reservoir storage is in the upper zone (greater than 400×10^6 m^3) withdrawals may be made for irrigation without limit, and for power, but not to exceed 250×10^6 m^3 per year for power.

3) When the reservoir is in the lower zone, the available water is apportioned between power and irrigation in the ratios of 20% to power and 80% to irrigation. This is the amount of water available for each use, not necessarily the amount actually used.

This agreement is not an operating rule, since it does not prescribe unambiguous releases. Rather, it is a constraint on the operation, favoring use of water for irrigation. Particular releases still must be determined in some manner, and should reflect this constraint.

Ibanez has developed a computer code that does not include the above constraints, but rather simulates a situation similar to the legal agreement, by the following proviso:

1) If the available water in Laguna Maule in any month is greater than 50% of the storage capacity, the excess over 50% must be released for irrigation downstream, except in February, when the entire active contents of the reservoir must be released downstream.

2) If the available water is less than 50% of capacity, then the release is zero. The operation of Invernada and Guaiquivillo is governmed by the standard operating policy and the associated

target values, and Colbun reservoir is treated as in the previous rule.

This policy was felt to be unsatisfactory, primarily because it does not reflect the constraining nature of the legal agreement, As a consequence, a policy termed the revised legal policy was developed by the author, to use the principles of the revised standard policy under the constraint of the legal agreement. Initially, the basic releases are determined, and are used to reduce the target values for the other two reservoirs in the upper basin. These two reservoirs are operated with a revised standard operating policy. The amount of water available downstream from the releases from Invernada and Guaiquivillo and the basic release from Maule is then determined, and compared with the irrigation requirements for the month.

If there is an irrigation deficit, and Maule is storing water in the upper zone, then water is used to supply this deficit from Maule. This water, supplied for irrigation, can also be used to satisfy the power requirements. Thus, it is compared with the target outputs for Maule reservoir, to see if a supplemental release for power is desired. If so, the supplemental release is made, with the provision that the total supplemental release in a year be less than 250×10^6 m^3.

If Maule is operating in the lower zone, the irrigation deficit is supplied from the 80% of available water allocated to irrigation, if possible. As before, the power releases may be made as supplemental releases to reach the target value, but not to exceed the 20% of available water allocated to power uses.

The complications of this policy are not as great as they might seem. The essential philosophy is to meet irrigation demands with the water allocated to irrigation. This water can also be used for power generation. A target release greater than the amount supplied to meet the irrigation deficit is taken to indicate a deficit in power production, and water is used from the "power pool" to supply this deficit, subject to the constraints on total and monthly percentage release for power.

4.4.5 Rule Curve Operating Policy

Ibanez developed an operating policy which assumes complete foreknowledge of all future inflows.for the period of simulation. The record is scanned for the dryest year in terms of total volume of flow. A rule curve which guarantees a firm yield based on the total flow of the dryest year is then developed by a fairly complicated procedure described in reference (15). A second rule curve is developed based on average flows. The operating policy then determines the available water in a given month. If this quantity is greater than the average for the month, then releases are made until the reservoir is at the appropriate point on the average rule curve. If the available water is less than the average for the month, releases are made to place the reservoir on the rule curve developed for the dryest year. Thus, the operation of the reservoirs is by attempting to maintain stored volume between the two rule curves, and where posssible, to minimize future spill in each reservoir by drawing down to the minimum volume in each month where this is possible and does not produce an expected deficit in future months. This procedure is carried out for each of the three reservoirs in the upper basin. Basic releases to avoid spills are computed, and the rule curve releases are added to the basic releases to obtain the total release. Colbun is again operated as before.

It should be noted that for this operating policy there is no interaction between the three upstream reservoirs. Each is operated independently according to its own rule curve. Further, no target releases are specified, so the policy is totally supply-oriented rather than demand-oriented.

4.4.6 Space Rule Policy

The space rule operating policy operates in a manner similar to the rule curve operating policy, generating minimum and average rule curves for reservoir storage, and calculating basic releases in each month. However, this rule provides some degree of interaction in the operation of the upstream reservoirs, as follows.

Instead of selecting the specific release within the constraining rule curves by the criterion of minimization of future spill for each individual reserovir, the operating policy responds to energy target values, and attempts to meet them wile being constrained to the rule curve, and operating so as to minimize total expected spill of the system over future months.

As before, a basic release is computed initially, and the energy derivable from the basic release is calculated and compared with the energy target for the current month. If an energy deficit exists, the three upstream reservoirs are examined for suplus water above the minimum rule curve. If only a single reservoir has water available, no allocation problem exists and this reservoir will either release to its minimum rule curve level, or release to just nullify the energy deficit, whichever is smaller. If more than one reservoir is available, a decision must be made on allocation of releases between the reservoirs. A complicated procedure attempts to satisfy energy deficits while using the space rule (8) to determine the allocation of releases in the upstream reservoirs that will minimize expected total spill of the system over the drawdown-refill cycle.

Colbun is again operated as a re-regulating reservoir for irrigation.

4.4.7 State-Oriented Policy

A policy was developed which provides some degree of state-orientation in the operation of the system. Essentially, the state-oriented policy revises the target values associated with the revised standard policy, in accordance with estimates of the available water for each reservoir for each month. The procedure is as follows:

1) Input to the state-oriented policy consists of monthly target drafts for each reservoir, monthly steady-state estimates of quantity of water available for each reservoir, and a parameter defining the shape of the function which modifies the targets. This input, with the exception of the parameter, is derived easily from the results of a linear programming model.

2) Then, in any month for any reservoir, the actual amount of water available, taken as the sum of storage plus inflow, is compared with the steady-state amount of water available, a program parameter. The concept of the operating policy is that if the actual water available is greater than the steady-state water available, the target release should be increased by some factor. Similarly, if the available water is less than the steady-state available water, the target should be reduced.

3) The determination of the amount of change in the target is made from the following function, where T is the original target value, T* the modified value, S the quantity of water actually available, W the expected quantity of water available, and p the value of the "shape" parameter, specified by the user:

$$T^* = T \times (p(S/W - 1) + 1)\ S/W \qquad (4.1)$$

A value of p of 0.0 provides fpr a direct linear proportionality between the modified target value and the ratio of water available to expected quantity of water available. A value of 1.0 makes the ratio of T^*/T proportional to $(S/W)^2$. Figure 4.4 is a plot of T^*/T vs. S/W for p as a parameter.

The philosophy behind this operating rule is that the linear programming model provides a steady-state estimate of available water for each reservoir in each month, and the linear programming targets are associated with this value. Therefore, a higher amount of water available should allow a greater draft without reducing future benefits significantly. This operating policy introduces 12 new parameters for each of the three upper basin reservoirs.in addition to the single shape parameter.

4.5 The Linear Programming Models

4.5.1 Introduction

During the course of the overall project research, a number of different linear programming models of the Maule basin were developed, based on varying data and assumptions, and with different numbers of periods per year. In the terminology of Chapter II, the Q set for each of these models varied in some of its elements.

The linear programming models were used to generate the target outputs required by the legal and standard operating policies, and the steady-state available water values needed for the state-oriented policy. Since the simulation is monthly, target outputs from 2-period and 4-period models were taken as uniformly spread over the appropriate number of months.

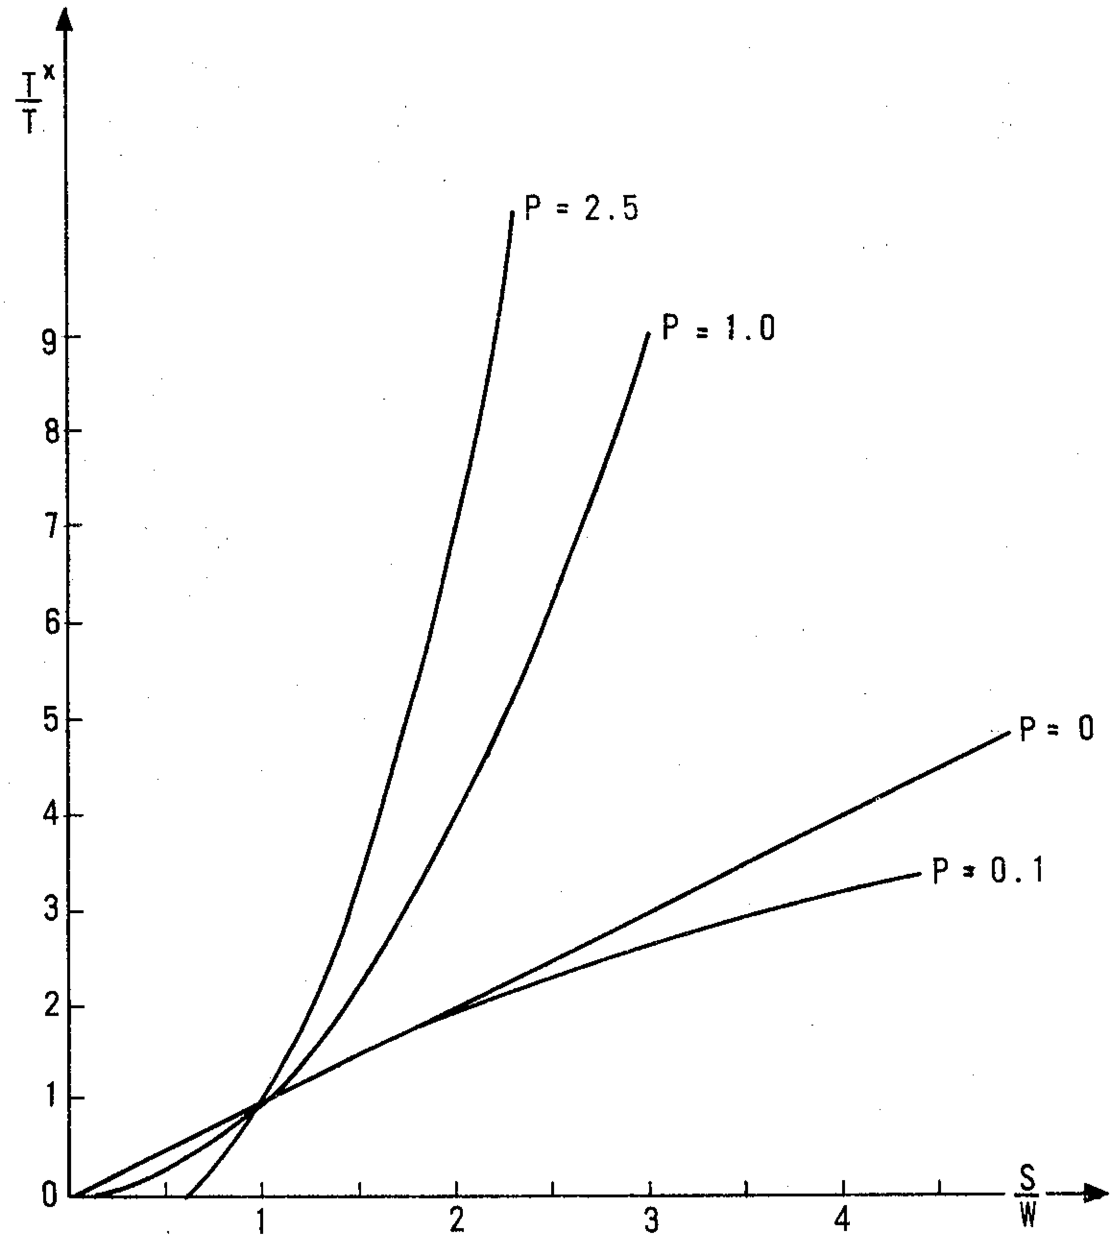

FIGURE 4.4

$\frac{T^x}{T}$ VS. $\frac{S}{W}$

4.5.2 Description of the Models

Ibanez used data from the 12-period model of the upper basin developed by Wallace in 1966 (5). This model was constructed at an early stage of the investigations into the Maule basin, and consequently the data, although the best that could be obtained at the time that the model was built, no longer represent the best information available. The three models developed by Poblete (20) are based on the most accurate data presently available. Data for the simulation model represents still another level, being developed intermediate to the work of Wallace and Poblete. Two sets of simulation data were used, the first being the originally developed data of Leonvendagar and McLaughlin, and the second, data upgraded in so far as possible by information obtained by Poblete.

Poblete has produced 2-period, 4-period, 12-period linear programming, and 4-period mixed integer linear programming models for the Maule basin. Results of the first three of these models were used in the simulation. Target values and available water quantities, as well as reservoir capcities obtained for Guaiquivillo and Colbun reservoirs, for each of the three models, are given in TAble 4.3. New information on reservoir capacity limitations for Guaiquivillo was obtained at a late stage of his investigation, and the values of Table 4.3 reflect the currently best data. Certain of Poblete's results prior to receipt of this new information were used with the original Leonvendagar and McLaughlin data for the standard and legal policies of Ibanez. Final results are obtained by a combination of the upgraded simulation data and the results of Poblete's investigations, including the new information on the reservoir capacity.

TABLE 4.3

RESULTS OF THE LINEAR PROGRAMMING MODELS

Month	2-Period Poblete			4-period Poblete			12-period Poblete			12-period Wallace			(Model)
	1	2	3	1	2	3	1	2	3	1	2	3	(Reservoir)
May	41.7	26.1	114.5	23.5	0.	103.1	0.	0.	49.3	111	4	299	Target Releases
June	↓	↓	↓	↓	↓	↓	66.9	0.	119.5	121	38	191	
July	↓	↓	↓	↓	↓	↓	38.7	0.	140.4	142	168	46	
Aug.	↓	↓	↓	27.9	0.	125.9	25.4	0.	126.7	127	71	171	
Sept.	↓	↓	↓	↓	↓	↓	27.2	0.	119.5	125	80	165	
Oct.	↓	↓	↓	↓	↓	↓	30.9	0.	131.5	127	76	220	
Nov.	65.7	42.1	172.0	69.1	0.	255.9	37.8	0.	228.2	201	0	164	
Dec.	↓	↓	↓	↓	↓	↓	74.8	0.	298.7	148	0	0	
Jan.	↓	↓	↓	↓	↓	↓	94.5	22.5	240.9	188	0	200	
Feb.	↓	↓	↓	94.5	146.6	88.0	94.5	79.5	128.0	157	81	120	
March	↓	↓	↓	↓	↓	↓	94.5	265.8	80.0	110	88	265	
April	↓	↓	↓	↓	↓	↓	59.4	71.1	55.9	107	197	263	

Month	2-Period Poblete			4-period Poblete			12-period Poblete			
	1	2	3	1	2	3	1	2	3	
May	41.7	30.5	114.5	23.5	30.4	103.1	62.1	22.6	49.3	Steady-State Available Water
June	↓	↓	↓	↓	↓	↓	80.4	54.8	119.5	
July	↓	↓	↓	↓	↓	↓	38.7	90.3	140.4	
Aug.	↓	↓	↓	27.9	61.1	165.9	25.4	122.6	126.7	
Sept.	↓	↓	↓	↓	↓	↓	27.2	152.5	119.5	
Oct.	↓	↓	↓	↓	↓	↓	30.9	182.4	131.5	
Nov.	83.2	47.1	172.0	125.7	121.9	255.9	85.3	230.9	228.2	
Dec.	↓	↓	↓	↓	↓	↓	208.2	305.3	298.7	
Jan.	↓	↓	↓	↓	↓	↓	264.5	365.0	240.9	
Feb.	↓	↓	↓	97.2	146.6	84.0	237.6	374.8	128.0	
March	↓	↓	↓	↓	↓	↓	176.6	318.4	80.0	
April	↓	↓	↓	↓	↓	↓	102.6	71.1	55.9	

CAPACITIES	2-Period Poblete	4-period Poblete	12-period Poblete
Gualquivillo	0.	0.	0.
Colbun	0.	577.5	698.2

All units - 10^6 m^3

4.6 Studies with the Model

4.6.1 Scope of the Studies

Studies with the Maule model principally involve search over the Q set involving different types of operating policies. All of the policies of Ibanez were tested and implemented for the first time, as was the model itself, since prior studies during the development phase were not produced. Thus, a major portion of the effort described in this chapter was involved with implementing and interfacing the model and operating rules.

Studies of Part I use the original data set of Leonvendagar, and the results of Wallace and early results of Poblete. For Part I, reservoir capacities for the proposed reservoirs were set at their maximum possible values, and the reservoirs were taken as initially full. Part II studes use the upgraded data set and the final form of the Poblete results. In accordance with the development of the previous chapter, reservoir capacities for proposed facilities were set at the levels derived from the linear programming results. The change in data for Guaiquivillo reservoir drove this reservoir to zero in all three of the linear programming models, and hence this reservoir does not appear in the Part II studies. The coding is such that target values and steady-state volumes of water for the non-existent reservoir must be suppdied to the simulation, and these are taken as run of the river values from linear programming, i.e., the target and the water available are identical and equal to the flow in the Melado river obtained from the linear programming model. Again, initially full reservoirs were examined.

Four synthetic hydrologies were used over a 40-year simulation duration.

Studies of Part I investigated interest rates of 2 through 16%, while the part II studies examined interest rate of 12, 14, and 16%, reflecting a typical xange of interest rates in the Chilean economy. Due to severe inflation, all economic data was normalized to give value in terms of escudos (E^o), the Chilean monetary unit, at their 1965 value (3.53 E^o of 1965 = $1 US), and benefit dat is presented as net benefit in millions of E^o of 1965.

Contrary to themmodels of the previous chapter, the simulation and the various linear programming models were developed independently, with some degree of interaction and joint development between Poblete and Leonvendagar and McLaughlin. There are, however, certain fundamental differences in the treatment of economic factors and the treatment of the entire irrigation sector between the two models, which prevent a direct comparison of the results in economic terms, The linear programming models of Poblete are felt to be sufficiently close to the simulation model to allow some degree of confidence in the levels of target and steady-state available water indicated by the linear programming model.

4,6,2 Part I studies

Part I studies investigate the relative merits of the various operating policies developed by Ibanez. The standard policy was studied with targets from the 2-period and 4-period models by Poblete, from the 12-period model of Wallace, and with target parameters prescribing "release" (empty reservoirs) and "store" (full reservoirs) extreme policies similar to those discussed in Chapter 3. The legal rule of Ibanez was tested with target parameters prescribed by the 4-period linear programming policy. Results for the different interest rates and the four hydrologies used are shown in Table 4.4 for all policies

TABLE 4.4

Benefit for Different Operating Policies

Hydrology	interest rate								
	2	4	6	8	10	12	14	16	Q
1	2196	1497	1015	679	439	264	132	31	
2	2400	1748	1295	975	744	573	443	342	RULE
3	2027	1424	1013	728	528	383	276	194	CURVE
4	1869	1264	849	561	358	211	103	22	
1	2327	1585	1075	721	468	284	146	41	
2	2543	1856	1380	1045	802	623	488	382	SPACE
3	2227	1570	1125	818	601	445	328	240	RULE
4	1995	1350	911	607	393	238	124	38	
1	2630	1810	1256	875	606	411	266	155	
2	2758	2011	1497	1137	879	689	546	436	LEGAL
3	2535	1808	1319	985	750	581	456	361	RULE
4	2351	1619	1127	791	556	388	265	173	
1	2599	1835	1313	951	695	510	372	267	
2	2783	2097	1619	1281	1037	857	721	616	Standard
3	2522	1845	1387	1072	851	693	577	489	(Wallace)
4	2358	1675	1209	888	662	501	383	294	
1	2573	1766	1221	845	579	386	243	133	
2	2731	1991	1482	1125	869	680	538	428	Standard
3	2475	1755	1271	940	708	541	418	324	(4-period)
4	2229	1528	1056	733	507	345	226	137	
1	2387	1642	1134	782	533	352	217	114	
2	2612	1916	1433	1091	845	663	525	419	Standard
3	2310	1643	1193	884	668	512	398	311	(2-period)
4	2088	1429	982	676	461	307	195	110	
1	-	-	-	-	-	-	-	-	
2	2793	2107	1628	1290	1046	866	720	625	Revised
3	2515	1838	1379	1064	843	685	569	482	Standard
4	2333	1652	1188	868	644	483	366	278	(Wallace)

(Benefit in escudos x 10^6)

except the extreme "release" and "store" policies. Mean inflows to each of the three reservoirs for the 40-year simulation are given in Appendix C for each hydrology.

Figure 4.5 presents results for the extreme policies. Net benefit is plotted against interest rate for both release and store policies, with hydrology as a parmeter. Figure 4.6 is a comparison of benefits for the three different standard policies obtained from the three different linear programming models, for hydrology 1.

4.6.3 Part II studies

Studies of Part II are made with the fully upgraded data set, that corresponds as well as possible to the data used by Poblete in his later models. The studies deal primarily with investigating changes in the Q set, to locate a suitable policy and suitable starting point. An example of a single-factor analysis on the operating policy parameters from such a starting point is also presented.

The investigations into the Q set are presented in Table 4.5, for four hydrologies, and interest rate of 12, 14, and 16% per year, for a forty-year simulation run. Table 4.5a presents results for the revised legal rule and for run of the river operation with zero reservoir capacities. The remaining tables that comprise Table 4.5 use either the revised standard or state-oriented policy.

The first column of the Q set in Table 4.5 provides information about the source of the policy parameter set. Values of 12, 4, and 2 refer to the parameter sets derived from the 12, 4, and 2-period linear programming models. A letter P indicates that the model used is that of Poblete, and a W indicates that the Wallace targets have been used. For runs 25 through 35, presented in Tables 4.5f and 4.5g, the target

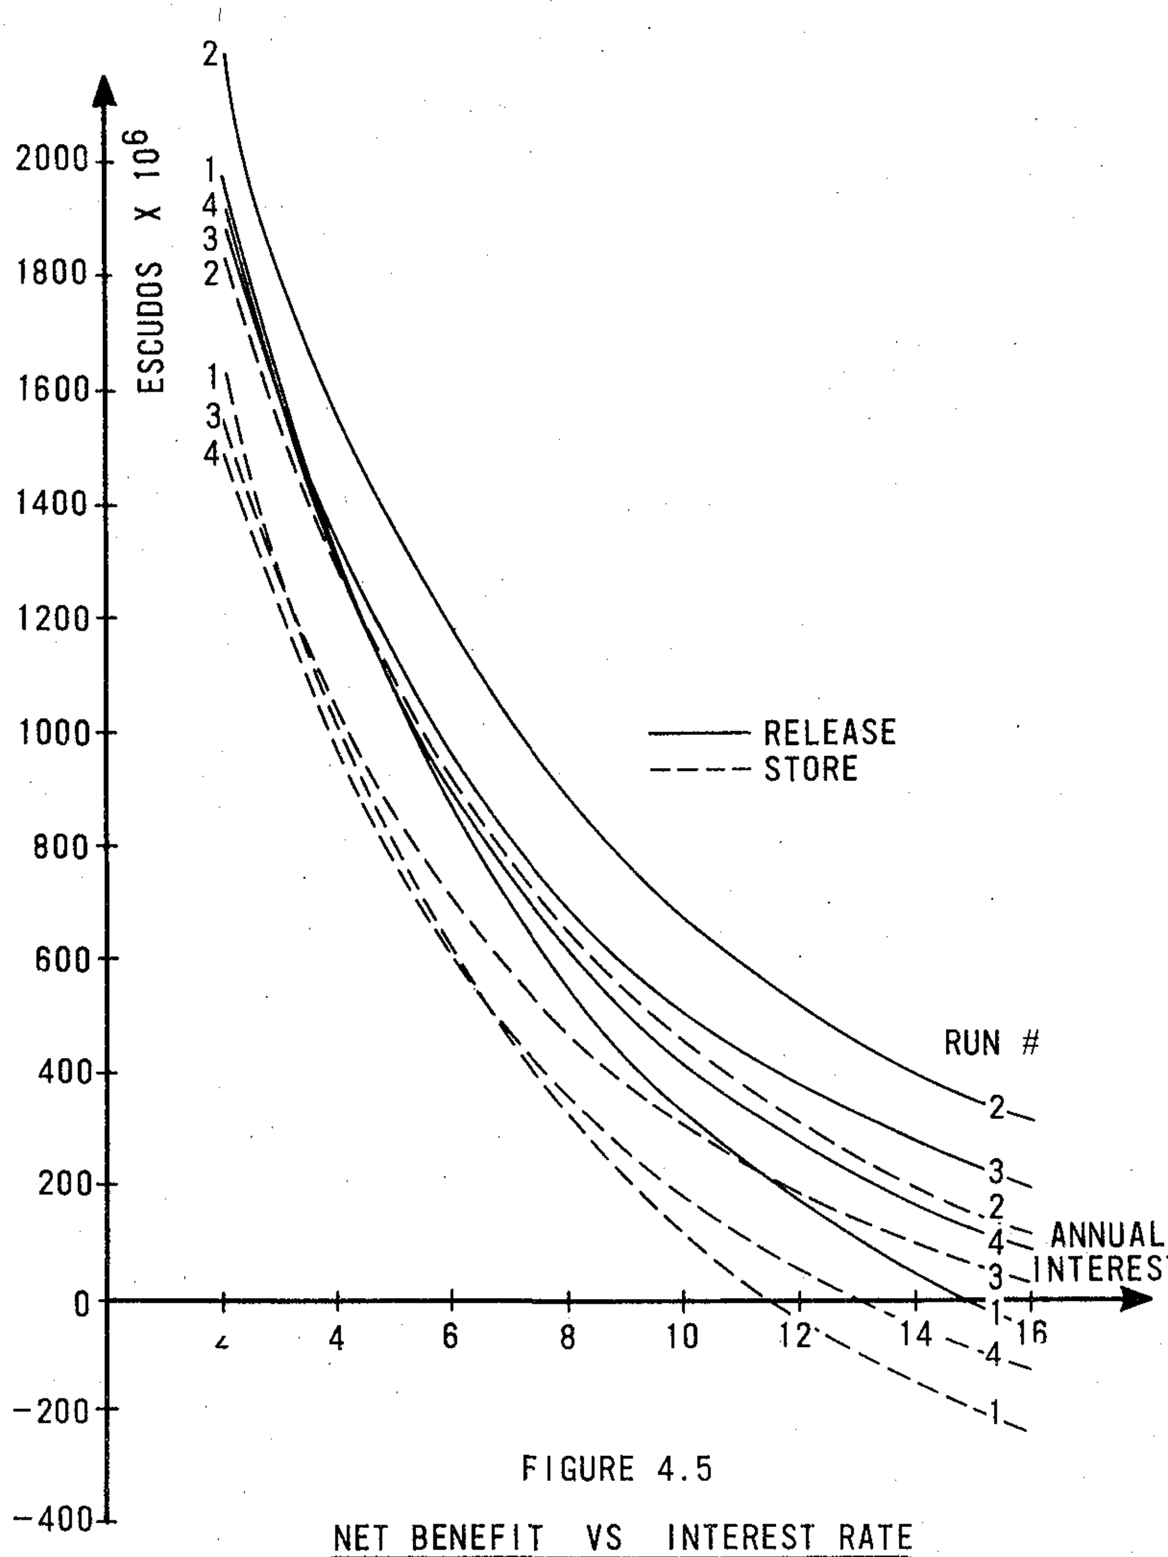

FIGURE 4.5

NET BENEFIT VS INTEREST RATE

RELEASE & STORE

STANDARD OPERATING POLICY

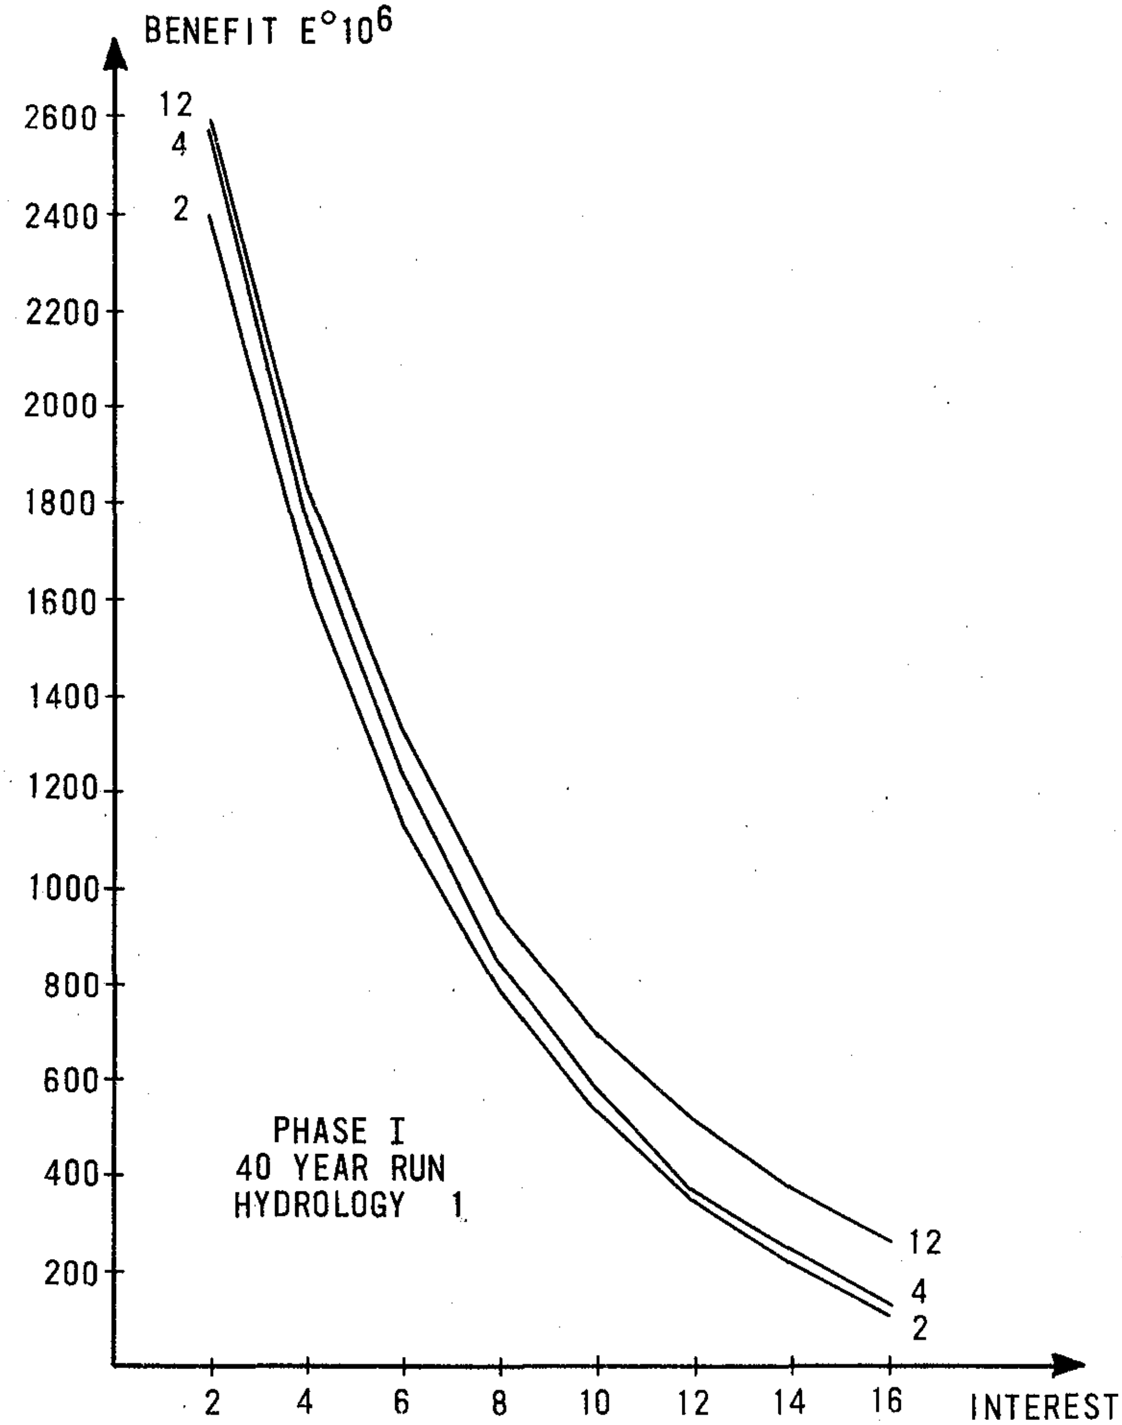

FIGURE 4.6

COMPARISON OF 3 LINEAR PROGRAMMING POLICIES

TABLE 4.5

MAULE STUDES - PART II

RUN	Hydrology	interest 12	14	16	Q
(1)	1	-132.9	-183.0	-224.0	CAPACITY = 0.
	2	159.1	110.8	72.9	
	3	43.7	11.6	- 13.2	
	4	- 57.9	- 94.5	-122.5	
(2)	1	178.2	57.9	- 35.0	
	2	455.4	340.5	252.6	LEGAL
	3	353.1	252.7	177.1	
	4	379.6	266.6	181.1	

(net benefits in millions of escudos)

Table 4.5a
Legal and Run of the River
Operation

RUN	Hydrology	12	14	16	Q		
(3)	1	276.0	160.4	70.6			
	2	572.2	459.5	372.7			
	3	454.5	358.1	284.8	12	RØ	
	4	419.9	312.3	230.1	P		
(4)	1	287.4	169.7	78.1			
	2	623.5	509.2	421.1	12	SØ	-.2
	3	496.8	397.9	322.7	P		
	4	462.9	352.0	267.2			
(5)	1	316.7	198.2	106.1			
	2	642.1	527.4	439.1	12	SØ	-.1
	3	512.3	413.3	338.1	P		
	4	489.0	378.5	294.2			
(6)	1	319.2	199.6	106.6			
	2	639.1	523.7	434.8	12	SØ	0.0
	3	514.3	414.6	338.8	P		
	4	485.6	374.7	290.1			
(7)	1	317.5	197.0	103.1			
	2	632.4	516.3	426.3	12	SØ	0.1
	3	507.2	406.7	330.2	P		
	4	-	-	-			
(8)	1	333.9	211.1	115.3			
	2	645.8	527.4	436.1	12	SØ	0.7
	3	522.9	420.1	341.9	P		
	4	486.6	373.1	286.3			

Table 4.5 b State and Release Oriented Policies

TABLE 4.5 (cont.)

Run	Hydrology	12	14	16	Q		
(9)	1	335.9	211.1	115.3			
	2	645.9	527.4	436.1			
	3	522.9	420.1	341.9	12	SØ	.8
	4	486.6	373.1	286.3	P		
(10)	1	336.2	212.9	116.8			
	2	648.2	529.2	437.5			
	3	526.6	423.1	344.3	12	SØ	1.0
	4	485.7	371.7	284.5	P		
(11)	1	334.4	211.0	114.8			
	2	645.3	526.3	434.5			
	3	527.6	423.9	345.0	12	SØ	1.1
	4	483.8	369.7	282.5	P		
(12)	1	312.2	189.1	93.2			
	2	641.9	522.9	431.2			
	3	528.1	424.2	345.1	12	SØ	1.2
	4	483.8	347.6	260.7	P		
(13)	1	341.8	217.3	120.4			
	2	641.2	521.5	429.4			
	3	532.0	427.1	347.8	12	SØ	2.5
	4	503.3	388.2	300.8	P		
(14)	1	340.6	215.8	118.7			
	2	644.5	524.4	431.8			
	3	533.3	428.0	348.0	12	SØ	3.0
	4	507.8	392.5	304.5	P		
(15)	1	321.4	196.0	98.4			
	2	628.9	509.5	417.4	12	SØ	5.0
	3	511.1	405.6	325.2	P		
	4	492.6	377.6	289.9			

(net benefits in millions of escudos)

Table 4.5b (cont.)

Run	Hydrology	12	14	16	Q		
(16)	1	332.6	213.6	121.0			
	2	616.9	502.6	414.3	4	RØ	
	3	508.6	409.9	334.7	P		
	4	461.6	351.9	268.1			
(17)	1	248.9	132.8	42.4			
	2	559.1	448.4	363.2			
	3	457.9	360.8	287.0	4	SØ	1.0
	4	379.9	274.7	194.3	P		
(18)	1	220.3	106.2	17.4	4	SØ	2.5
	2	531.3	422.7	338.9	P		

($E^o \times 10^6$) Table 4.c - 4-period Poblete targets

TABLE 4.5 (cont.)

Run	Hydrology	12	14	16	Q		
(19)	1	37.7	-27.1	-78.2			
	2	291.8	224.4	171.6	2	RØ	
	3	173.7	128.2	93.0	P		
	4	183.3	120.6	72.0			
(20)	1	6.6	-56.5	-106.4			
	2	228.8	167.7	120.4	2	SØ	0.
	3	103.0	61.3	29.2	P		
	4	116.5	60.0	16.3			
(21)	1	-38.4	-100.2	-149.3			
	2	177.3	121.2	77.4	2	SØ	1.0
	3	67.9	30.2	1.0	P		
	4	77.1	23.9	-17.2			
(22)	1	-86.7	-145.1	-191.7			
	2	166.9	113.4	71.2	2	SØ	2.5
	3	38.6	3.7	-23.2	P		
	4	27.8	-20.8	-58.2			

(net benefit in escudos x 10^6)

Table 4.5d
2-Period Poblete Targets

Run	Hydrology	12	14	16	Q		
(23)	1	359.2	244.9	156.3			
	2	676.7	565.3	479.6			
	3	537.7	441.3	368.2	12	RØ	
	4	510.0	402.7	321.0	W		
(24)	1	338.1	214.4	118.4			
	2	671.2	551.5	458.9			
	3	547.7	441.4	360.3	12	SØ	1.0
	4	511.4	396.6	308.9	W		

(net benefits in escudos x 10^6)

Table 4.5e
12-period Wallace targets

TABLE 4.5 (cont.)

Run	Hydrology	12	14	16	Q	
(25)	1	278.7	165.2	77.0		
	2	586.1	477.6	394.0	1.2x	RØ
	3	468.3	373.8	302.0	12 p	
	4	433.8	327.3	246.0		
(26)	1	287.8	173.9	85.3		
	2	579.0	471.5	388.8	1.6x	RØ
	3	466.4	372.3	301.0	12 p	
	4	435.4	329.9	249.3		
(27)	1	293.4	178.9	89.8		
	2	587.5	479.9	397.2	1.8x	RØ
	3	470.9	377.1	306.0	12 p	
	4	447.0	341.2	260.3		
(28)	1	294.8	179.7	90.1		
	2	585.3	477.7	394.9	2 x	RØ
	3	471.5	376.8	305.0	12 p	
	4	458.5	352.6	271.7		
(29)	1	289.4	173.1	82.5		
	2	563.3	455.2	372.0	4 x	RØ
	3	463.6	368.1	295.5	12 p	
	4	427.3	320.9	239.6		
(30)	1	282.9	167.6	77.7		
	2	558.6	451.4	368.8	6 x	RØ
	3	449.0	354.6	283.0	12 p	
	4	422.2	316.8	236.2		

(net benefit in escudos x 10^6)

Table 4.5f. Sensitivity to Target Magnitude

TABLE 4.5 (cont.)

Run	Hydrology	12	14	16	Q		
(31)	1	318.2	194.3	48.0	1.2x 12 P	SØ	3.0
	2	628.1	509.2	417.5			
	3	521.3	416.5	336.9			
	4	498.4	384.3	297.2			
(32)	1	303.8	182.5	388.4	1.6x 12 P	SØ	2.5
	2	620.8	504.3	414.8			
	3	511.7	408.0	329.4			
	4	477.5	365.5	280.2			
(33)	1	288.1	168.2	75.1	1.8x 12 P	SØ	3.0
	2	599.5	484.4	396.0			
	3	501.5	398.8	320.9			
	4	455.6	354.4	261.5			
(34)	1	192.6	78.6	-9.8	4 x 12 P	SØ	3.0
	2	537.8	426.3	340.8			
	3	441.9	344.5	270.7			
	4	401.5	295.1	214.0			
(35)	1	298.3	175.6	80.0	2 x 12 P	SØ	2.5
	2	600.3	484.9	396.0			
	3	475.8	373.7	296.0			
	4	440.7	328.7	243.3			

Table 4.5g, Sensitivity to Target Magnitude, State-oriented

TABLE 4.6

SINGLE FACTOR ANALYSIS

Hydrology	Target	Change in target	12	14	16	difference in benefit
1	RØ		359.2	244.9	156.3	
2	12 - Wallace		676.7	565.3	479.6	
3	starting		537.7	441.3	368.2	
4	point		510.0	402.7	321.0	
1	(1,1)	+10	358.9	244.6	156.0	- .3
3	(1,2)		537.7	441.3	368.2	=
1	(1,3)		359.2	244.9	156.3	=
3	(1,4)		537.7	441.3	368.2	=
4	(1,5)		510.2	402.9	321.1	+ .2
1	(1,6)		359.2	244.9	156.3	=
3	(1,7)		537.7	441.3	368.2	=
2	(1,8)		675.7	564.3	478.6	- 1.0
1	(1,9)		359.0	244.7	156.1	- .2
2	(1,10)		675.3	563.7	478.3	- 1.4
3	(1,11)		537.7	441.3	368.2	=
4	(1,12)		510.1	402.8	321.1	+ .1
2	(2,1)		677.0	565.7	480.0	+ .3
4	(2,2)		510.4	403.1	321.4	+ .4
2	(2,3)		675.7	564.3	478.6	- 1.0
4	(2,4)		510.2	403.0	321.2	+ .2
3	(2,5)		537.7	441.3	368.2	=
2	(2,6)		676.3	564.3	479.2	- .2
4	(2,7)		507.4	400.5	319.1	- 2.6
1	(2,8)		356.0	242.2	153.9	- 3.2
1	(2,9)		356.9	243.0	154.7	- 2.3
2	(2,10)		673.7	562.8	477.4	-3.0
3	(2,11)		536.0	440.0	366.9	- 1.7
4	(2,12)		510.4	403.0	321.2	+ .4

(escudos x 10^6)

TABLE 4.6 (cont.)

Hydrology	Target	Change in target	12	14	16	difference in benefit
2	(1,1)	-10	676.7	565.4	480.0	=
3	(1,2)	↓	537.7	441.3	368.2	=
4	(1,3)		510.0	402.7	321.0	=
1	(1,4)		359.2	244.9	156.3	=
2	(1,5)		676.7	565.3	479.6	=
3	(1,6)		537.7	441.3	368.2	=
4	(1,7)		510.0	402.7	321.0	=
1	(1,8)		359.5	245.1	156.5	+ .3
2	(1,9)		678.2	566.4	481.1	+ 1.5
3	(1,10)		538.4	441.6	368.4	+ .7
3	(1,11)		537.9	441.5	368.4	+ .2
4	(1,12)		510.1	402.8	321.0	+ ..1
1	(2,1)		359.0	244.7	156.0	- ..2
2	(2,2)		676.5	565.1	479.3	- ..2
3	(2,3)		539.9	443.5	370.3	+ 2.2
4	(2,4)		509.8	402.5	320.7	- .2
1	(2,5)		358.7	244.3	155.7	- .5
2	(2,6)		676.9	565.5	479.8	+ .2
-	(2,7)		-	-	-	
-	(2,8)		-	-	-	
-	(2,9)		-	-	-	
3	(2,10)		541.2	444.5	371.0	+ 3.5
4	(2,11)		510.1	402.7	321.0	=
1	(2,12)	↓	359.3	244.9	156.4	+ .1
1	Revised Target Values RØ^		358.8	244.2	155.6	
2			687,2	574.9	488.4	
3			544.1	447.1	373.5	
4			514.4	406.4	324.1	

(escudos x 10^6)

Revised Target Values (10^6 m^3)

month	Res. 1	Res.2
1	101	4
2	121	48
3	142	158
4	127	81
5	125	90
6	137	66
7	201	0
8	138	0
9	178	0
10	147	71
11	100	78
12	117	187

values have been increased by multiplication by the amount noted. For example, in run 25, the 12-period Poblete targets have been increased by a multiplication factor of 1.2.

The second column of the Q set provides a sumbol indicating whether the policy is state or release oriented. If the symbol SØ appears in the second column of the Q set, the policy is the state-oriented policy described in Section 4.4.7. In this case, the third column of the Q set gives the value of the parameter p used in association with the policy. If the symbol RØ appears in the second column, the policy is the release-oriented revised standard policy described in Section 4.4.3. For this case, the third column is blank.

Table 4.5b presents studies made with the 12-period Poblete model to compare the release-oriented rule with the state-oriented rule for different values of the "shape" parameter p.

Table 4.5c presents results for simulation runs using the 4-period Poblete model parameter set. Benefits are presented for the release-oriented policy and for the state-oriented policy with two different values of the shape parameter.

Net benefits for the release-oriented policy and three different state-oriented policies are presented in Table 4.5d for the 2-period model.

The 12-period Wallace targets are used to provide operating policy parameters for runs 23 and 24, presented in Table 4.5e. For the state-oriented policy, the 12-period Wallace release targets are used with the 12-period Poblete values for steady-state available water.

Studies presented in Table 4.5f investigate the sensitivity of benefits to the magnitude of the policy parameters for a release-oriented rule. The 12-period Poblete targets are multiplied by the indicated factors, thus preserving the same ratio of release targets between months and reservoirs, but changing the absolute level of these releases.

Table 4.5e continues the studies of sensitivity to magnitude of the policy parameters for state-oriented policies. For runs 31 through 34, both the release targets and the values for water available obtained from the 12-period Poblete model are multiplied by the indicated factors. For run 35, only the release targets are multiplied by a factor of 2, while the estimates of available water are held at the original level obtained from the 12-period Poblete model.

Table 4.6 presents an example of single-factor analysis, starting from a release-oriented policy using the 12-period Wallace targets. Since the results of the linear programming models show that Guaiquivillo reservoir should not be constructed, the target releases for Guaiquivillo need not be examined. Consequently, the single-factor analysis was performed over the 24 different target values, one for each month for both reservoirs 1 and 2. The grid spacing was arbitrarily selected as 10 million m^3, and the targets were varied in each direction where possible. Targets (2,7) (reservoir 2, month 7) through (2,9) were at zero level, and thus could not be diminished. The hydrology used varied among the four hydrologies used in the study, and the particular hydrology used for each parameter change is indicated, together with the change in benefits from the original target values. A new starting point derived from the information obtained by the single-factor analysis is presented, together with the associated benefits for all hydrologies.

4.7 Discussion of Results

4.7.1 Part I

Results presented in Table 4.4 indicate that maximum benefits are derived from the standard policy with the 12-period Wallace targets, except for the 2% interest value for hydrology 1, for which the legal policy proves slightly superior. The space rule and rule curve policies yield benefits significantly lower than the legal or standard policies, and as a consequence are not considered in the Part II studies, but note should be taken of the fact that introduction of some degree of optimization in the space rule policy gives this policy an advantage over the rule curve policy in terms of benefit levels generated.

Data from Table 4.4 for the three different linear programming models used is presented graphically in Figure 4.6. Examination of this figure shows a fairly constant difference in benefits between the 12 and 2 period model results. The 4-period model at low interest rates is nearly as good as the 12 period model, while for the higher interest rates, it is only slightly better than the 2-period model. The linear programming models were constructed with an interest rate of 15%, for which the 2 and 4 period models yield fairly close benefits in the simulation.

The behavior of the three different models for different interest rate can be explained as follows. For a low interest rate, all monthly benefits are weighted close to equally in the discounting process, while for the higher interest rates, there is a rapid decrease in the effect of benefits accrued in later stages of the imulation. Thus, at lower interest rates, the 4 period model, by virtue of its greater accuracy in reflecting the monthly changes, yields values close to the 12-period model results. For the higher interest rates, the greater accuracy of the 12 period model becomes apparent as the proper sequencing of releases in the early stages of the simulation becomes more critical. In the range of practical interest rates (12-16%), there is a clear advantage to using the results of the 12-period model.

Results of simulation runs with the revised standard policy are presented in Table 4.4 for hydrologies 2, 3, and 4. Results for hydrology 2 show an increase in benefits over the standard policy, while benefits decrease compared to the standard policy for the other two hydrologies. Reference to Appendix C will show that hydrology 2 is the "wettest" of the 4 hydrologies used, with largest volume of total inflow. The behavior of the simulation with revised standard policy can be explained as follows.

The revised policy allows the system to end a month with a full reservoir, which is only possible in the standard policy if the target draft is zero. The higher level of reservoir in the revised policy will create larger basic releases in a wet year. For particularly large inflows, the basic release may be much greater than the controlled releases, while still maintaining a full reservoir. Benefits are obtained from the basic releases while water is conserved in the reservoir for later use in periods of scarcity. Therefore, in a wet year the revised policy will yield higher benefits due to a greater supply of water. The standard policy will draw down reservoirs by wasting water in periods of abundance, when the basic releases satisfy target demands without supplemental releases. The standard policy makes the supplemental releases in any case, wasting this water.

For detailed comparison between the revised and standard operating policies, comparison of the physical operation of the system for the two different policies should be performed.

Figure 4.5 presents a study of the two extreme policies of "release" and "store". The rank ordering of the two policies is seen to change, both with interest rate and as a function of the particular hydrology used.

Comparison of these two policies allows an examination of the effect of initial volume.

For all runs with the Maule model, all reservoirs were taken as initially full to capacity. Thus, the release policy will immediately supply the initial volume, while the storage policy will immediately operate as a run of the river policy, since the reservoirs are already full to capacity. The difference in benefits due to initial volume can be determined by an examination of the release and store policy for the same hydrology. This difference is clearly a function of interest rate. Benefits obtained from supplying the initial volume by a release policy at the start of the simulation will be weighted proportionately more in the discounted evaluation for higher interest rate. Results derived from Figure 4.5 confirm this. For an interest rate of 16%, the difference in benefits between the two policies for hydrology 2 is seen to be 191.6 million escudos, while for a rate of 2% the difference is 339.5 million escudos.

Release and store policies are interesting for the fact that they can isolate the effects of initial volume, and because the manner of operation they prescribe is easily comprehended. Certain situations of planning and design may actually find that these policies are limiting cases. For a variable head power plant, operation may be optimal for the storage policy, which maintains high head at the turbines. In certain over-year storage situations, operation for a large portion of the year may be run of the river with empty reservoirs.

4.7.2 Part II

Table 4.5a presents results for run of the river operation of the basin with zero capacities and for operation with the revised legal rule. Large negative net benefits are indicated for uncontrolled flow for all hydrologies except the "wet" hydrology 2.

Examination of detailed output for results using the revised legal rule indicates that the rule does indeed operate to favor agricultural production. For hydrology 1 and 12% interest rate, net agricultural benefits are 116.6 million escudos higher than net agricultural benefits for the release-oriented policy using the 12-period Wallace targets. This increase is obtained at the expense of a decrease of 297.2 million escudos in net benefit for power production from the value obtained by the 12-period Wallace model.. This indicates that the subsidy to agriculture implied in the legal agreement is a costly one.

Table 4.5b presents results of investigations into the shape of the state-oriented policy using the 12-period Poblete targets. Shape of the policy is determined by the shape parameter p. Referring to Figure 4.4, it can be seen that all curves of ratio of revised target draft to draft vs. ratio of available water to steady-state available water pass through the point (1,1). This implies that for all shapes of the policy, if the actual water available in any month is equal to the steady-state value derived from the linear programming model, then the target draft should also equal that derived from the linear programming model.

A shape parameter value of 0.0 provides a linear relationship between the two ratios, while a value of 1.0 provides a quadratic relationship. Certain of the curves shown in Figure 4.4 intersect the axis, and derived target drafts are negative. Whenever this is the case, a provision of the operating rule sets target drafts to zero. Thus, if the water available is sufficiently less than estimated, there will be no controlled release. For increasing values of p beyond 1, the policy becomes increasingly more conservative of deficient water as the point of intersection moves along the axis. For increasing values of p, the policy also becomes increasingly more extravagant of excess water as the shape curve becomes steeper. It must be noted that an increase in the target to 9 times the linear programming target release does not necessarily mean that 9 times the quantity of water will be supplied. Actual quantity of water available in a period limits the total release. For a sufficiently large value of the shape parameter, the policy becomes essentially bi-stable, releasing entire available contents when hte available water is just slightly greater than estimated, and releasing zero when the available water is slightly less than estimated.

Results presented in Table 4.5b show that the state-oriented policy is superior to the release oriented policy in all cases. A shape parameter of 1.0 yields maximum benefits for runs using hydrology 2, 3.0 is superior for hydrologies 3 and 4, while a value of 2.5 is obtained for hydrology 1. Benefits generally follow a smooth curve with the value of p, but there are some deviations. This indicates that a small change in the target releases may prove critical in a particular hydrology and produce a large effect on benefits due to the cumulative effects of a decision in a state-oriented rule. A particular example is the unexpectedly low benefit obtained for run 12, using a parameter value of 1.2, for hydrologies 1,2, and 3. Detailed examination of the physical output would be required to explain precisely how the change in the value of the shape parameter causes this decrease in benefits.

Results for simulation runs using target parameters derived from the 4-period Poblete model are shown in Table 4.5c. Results show two interesting features. The 4-period release-oriented policy proves superior to the 12-period release-oriented policy, and benefits decrease in the transition to the state-oriented policy. The 4-period release-oriented policy does prove inferior to the best 12-period state-oriented policy however,

This behavior can be explained by postulating that the nature of the 12-period model is such that target values are sensitive to estimates of water available, while for the 4 period model, estimates of water available are not accurate. The state-oriented 4-period policy, in revising target drafts based on estimates of steady-state values, is revising on the basis of inaccurate information, and consequently leads to lower benefits than the unrevised target drafts of the release-oriented policy. Similarly, the greater sensitivity of the 12-period model implies that the target drafts are closely tied to the steady-state available water values. Use of these target drafts without reference to actual state in the release-oriented policy is an inappropriate use of these values, as they are conditioned by the state values, and therefore lead to lower benefits than the 4-period release-oriented policy.

The above is not a definitive explanation of the behavior of the two models, and further study of such behavior is indicated.

Results for the 2-period Poblete targets are presented in Table 4.5d. Again, benefits decrease as state-orientation is introduced, but the 2-period release-oriented policy is far inferior to either the 4- or 12-period policies, indicating that the two-period model is inadequate for predicting appropriate target releases for a 12-period simulation.

Table 4.5e shows results for operating policies using the 12-period Wallace targets. The Wallace linear programming model does not correspond to the present simulation model in a number of features, notably that much larger irrigated farming areas are feasible in the Wallace model. As a consequence, the target releases from the Wallace model are significantly larger than those of the Poblete model.

The state-oriented policy for the Wallace targets was obtained by using the 12-period Poblete estimates of water available. Results indicate that for hydrologies 3 and 4 at an interest rate of 12%, there is improvement in benefits over the release-oriented policy, while results for the remainder of the studies show a decrease in benefits as state-orientation is introduced. No sensitivity analysis to the value of p was performed for this case, and the value of 1.0 was selected before the studies with the Poblete model indicated that other values of the parameter might be appropriate. It is expected that a sensitivity analysis on this parameter would improve benefits of the state-oriented policy relative to the release-oriented policy. As noted above, however, the 12-period model with a state-oriented rule is sensitive to the combination of the target and estimated state parameters. The fact that the Poblete steady-state values are used with the Wallace target parameters indicates that a more appropriate set of state parameters for the Wallace targets may exist which would yield obvious improvement of the state-oriented policy over the release-priented policy. The release-oriented policy for the 12-period Wallace targets yields the highest benefits found during the studies of Table 4.5 with the exception of the improvement obtained by moving to a state-oriented policy for 12% interest rate and hydrologies 3 and 4.

Table 4.5f presents results of studies using a release-oriented policy to determine sensitivity of benefits to the magnitude of target releases. To reduce the number of parameters that would be involved by individually varying each target, the entire 12-period Poblete target set was multiplied by a variable parameter. Depending on the hydrology chosen benefits generally increase for values up to between 1.8 and 2.0 times the linear programming values, with increases in benefit over the results using the original target values on the order of 10 to 15 million escudos. Too great an increase in the target values causes a decrease in benefits due to wastage of water beyond the possibilities for benefical use.

Similar studies are performed for the state-oriented policy, and results are presented in Table 4.5g. For all runs except run 35, both the target and estimated state parameters from the 12-period Poblete model are multiplied by the factor indicated. For the final run, only the target release parameters are multiplied by two. Results show an interesting trend. For multiplication factors of 1.2 and 1.4, the state-oriented policy shows a fairly large increase in benefits over the release-oriented policy, but for increasing values of the multiplication factor, benefits decrease sharply. This indicates once again the sensitivity of the state-oriented policy to the estimated state parameters. Bad estimates produce significant decreases in benefit.

The final run presented in Table 4.5g, revising only the target parameters, has reversed the downward trend in benefits with increasing values of the multiplication factor, and suggests revising only the target parameters in a state-oriented policy is a more appropriate sensitivity study.

Such a study would take into account the fact that doubling the target volume is likely to change the quantity of water available in each month, but it is not likely to double it. The above studies are run on the basis of a doubling in scale of the entire linear programming model, which is clearly not a valid approach. A preferred approach would either use a statistical study of physical operation or an adaptive control process to revise state parameters to appropriate values for different magnitudes of target values. An alternate approach, somewhat simpler in practice, would be to re-solve the linear programming problem with the given target levels as constraints rather than decision variables. The linear programming equations would then yield the appropriate values for water available.

Further search on the form of the operating rule was not pursued beyond the 35 runs presented in Table 4.5, with the exception of a small-scale test of a state-oriented rule which revised target releases based on the difference rather than the ratio between the actual and steady-state water available. Results are not presented here since benefits generated were consistently lower than benefits from the present state-oriented rule.

It is obvious that there are an infinite variety of possible operating rules, and even those of simple form can be developed at a rapid rate with a small amount of effort. At this point, the LPSS process could proceed in the direction of further improvements in the Q set, or can continue the search on the values of the operating policy parameters, as done in Chapter 3. It was recognized that improvement in the Q set should take the direction of development of an appropriate adaptive control rule based on values from linear programming, and characterized by as few parameters as possible.

Rather than proceed in this direction, which would have been a major involvement in an investigation on the borderlines of the main area of concentration of this thesis, a single-factor analysis on the parameter set of target releases from the 12-period Wallace model, generally yielding highest benefits, was performed.

Results, as presented in Table 4.6, show that only a small number of variables produce significant changes in benefit. The single factor analysis required 45 computer simulation runs, each requiring approximately 1.2 minutes of computer time on the IBM 360/65, to yield a parameter set which decreased benefits slightly for hydrology 1, increased benefits for hydrology 2 by 10 million escudos, increased benefits for hydrology 3 by 7 million escudos, and for hydrology 4 by 4 million escudos. The incremental benefit increase due to this portion of the search process was small and computer time and data handling effort were large compared to the previous search on the Q set. Consequently, the LPSS process was terminated at this point.

4.7.3 General discussion of Results

Portions of the LPSS process carried out for the Maule model are presented in Figure 4.7, with benefit presented as a function of the "search process". As noted previously, neither the linear programming model of Wallace nor Poblete corresponds precisely with the simulation. Each model may be thought of as an estimate of a hypothetical "best" linear programming model for the simulation. It remains for further studies to determine criteria for construction of such a "best" linear programming model for a given simulation.

The LPSS process should eventually converge to a near-optimal policy independent of the linear programming model used. The rapidity of this convergence will clearly depend upon the initial starting values determined by linear programming. For the Maule basin studies, the Wallace model yields a superior initial value to the Poblete model, and consequently the LPSS process starting from this model is shorter and yields a higher benefit than the search starting from the Pobdete

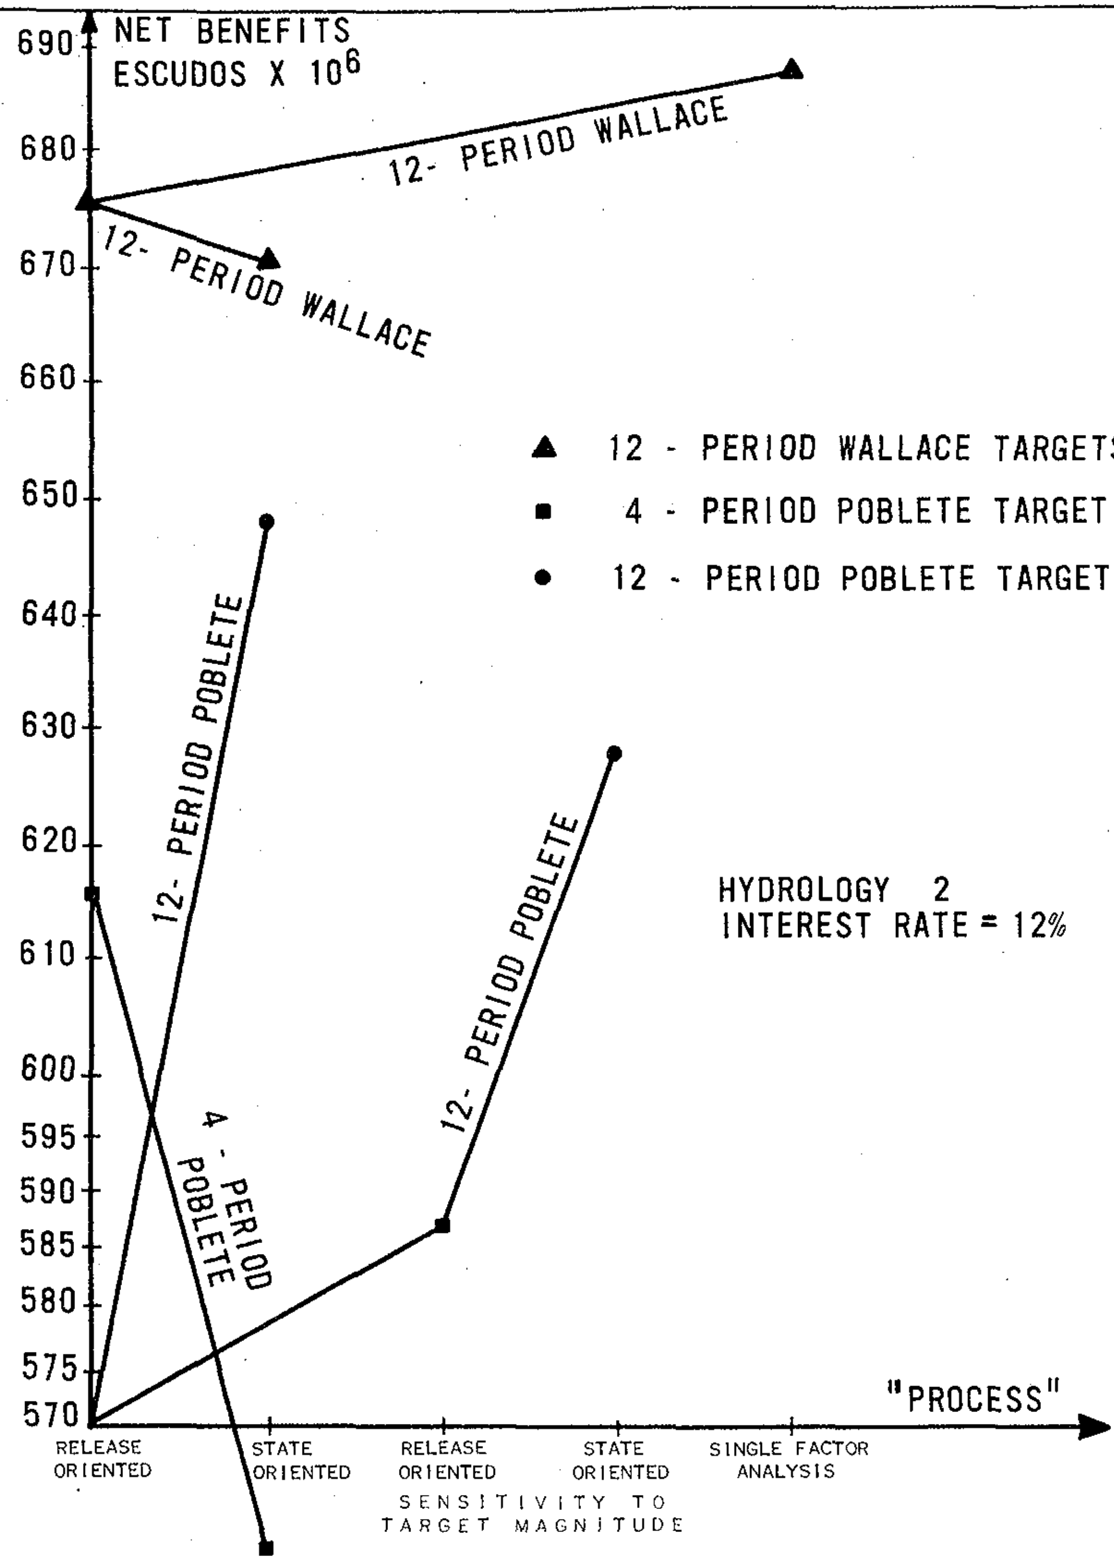

FIGURE 4.7
LPSS PROCESS CONTROL

model. It should be noted that the method of plotting Figure 4.7 tends to exaggerate differences, since the origin is not taken at 0.0 benefits but at 570 million escudos. Percentage increases in net benefits are on the order of 10 to 20%, depending on the particular stage of LPSS.

Figure 4.7 shows that two different search paths were taken using the 12-period Poblete model. The first path simply used the maximum amount of data from the linear programming model, and performed a search over the Q set to differentiate between state- and release-oriented policies, and to define the "shape" of the state-oriented policy. The other search path examined response to magnitude levels of the parameters of the operating policy. Both search paths yield significant increases in benefits. These search paths were terminated due to recognition of the problem of obtaining steady-state parameters which correspond to target parameters, to be discussed below.

For the state-oriented policy, there are 24 operating policy parameters for each reservoir, 12 target releases and 12 steady-state available water values. Clearly there is a relationship between the values of target releases and the values of steady-state water available. For the state-oriented policy to be an improvement over the release-oriented policy, the steady-state available water parameters must be consistent with both the target parameters and "reality" for the simulation model. As has been seen, inconsistency in either of these manners yields decreases in benefits as the state-orientation has been introduced. Consequently, when LPSS is pursuing search on the parameters of the operating policy, a consistent set of the 24 parameters must be obtained for each reservoir. Such consistency can be obtained by searching over all the parameters, but the large number of variables makes this expensive. A method is needed to provide consistent state parameters while the search process is carried out on the release parameters. In the absence of such a method, the search paths from the 12-period model are terminated at the point at which detailed search of the target parameters ordinarily would begin. Suggested methods for developing consistent parameter sets are given in Chapter 6.

As a consequence of the foregoing considerations, the search path using the 12-period Wallace targets was based on the release-oriented policy. Only the single-factor search was performed. Benefits were increased by 1.6% due to performance of the single-factor analysis, for hydrology 2. Termination of the search path stemmed from two considerations. First, results of Chapter 3 indicated that incremental benefit due to search after single-factor analysis decreases rapidly. Secondly, the data structure for the Maule model proved very difficult to handle with respect to searches over the policy parameters. Additional search effort is large for the Maule model, and consideration of the possible advantages of continuing search by another means led to the decision to terminate the search path at the point shown. In the absence of methods of improving the form of the policy (e.g. developing a consistent state-oriented policy), the current policy was accepted as sufficiently near-optimal.

For purposes of research into the LPSS process, the Maule model was treated as a black box. Input was the Q set, and output was the single value of net benefits, for each interest rate. No detailed examination of the physical operation of the basin was performed, and the search was conditioned solely by the benefits obtained. For implementation to determine near-optimal policies for a given simulation model however, consideration of physical outputs should aid in convergence, as it did in Chapter 3 studies.

Results for the Poblete model show that for a release-oriented policy, the model underestimates values of the target parameters. This consistent characteristic of the linear programming model can be understood as follows.

Linear programming target values arise as releases which maximizes the objective function, subject to constraints on available water and capacity. The capacity constraint will limit target releases to values which do not exceed maximum downstream useful flow, provided that reservoir storage capacity is not exceeded. The available water as determined by the steady-stae hydrology also serves to constrain the targets. Thus, if steady-state available water is augmented, the target releases can be increased if capacity constraints downstream are not active.

Stochasticity in the simulation will cause inflows and quantities of water available which are occasionally greater than and occasionally less than the steady-state linear programming values. Target releases determined by linear programming models for which the available water constraint is active will limit the maximum release, and target releases may not be met. Consequently, the setting of higher target values in the simulation model than those determined by the steady-state linear programming model will not cause loss of benefits in periods of low flow, since available water controls, and will allow for increased benefits in periods of abundant water, by releasing excess water for beneficial use.

The Maule studies clearly showed the superiority of the 12-period models over the 4 and 2-period models, particularly when state-orientation is introduced. The assumptions required to transfer results of a 2 or 4 period linear programming model to parameters for a 12-period model do not appear to be valid. It cannot be concluded that the 12-period linear programming model is superior to the 4-period model for use in simulation studies. Rather, it can be noted that for a 12-period simulation model, a 12-period linear programming model is appropriate to yield target parameters. The 4-period model may prove entirely adequate when used in conjunction with a 4-period seasonal simulation, but at present it does not appear possible to "force" the 2 and 4 period linear programming models to obtain appropriate target values for a monthly simulation model.

General experience with LPSS for the Maule model indicated that data handling and data structure were inefficient. Search techniques require a relatively large number of relatively small modifications, often to only a single parameter at a time. For the current data structure, excessive amounts of time are required to organize punched-card data for the search process. Therefore, the existence of the LPSS concepts points out that data structures for simulation models should be organized in a manner that makes search efficient. A suggested data handling method will be described in Chapter 6.

4.8 Conclusions

Principal conclusions from the Maule study are the following:

1) The LPSS process, treating the simulation as a black box, does show significant increases in benefit, although convergence is not as rapid or obvious as that obtained in Chapter III.

2) The introduction of state parameters to obtain a state-oriented rule requires that the state parameters be consistent with the target releases.

3) Since benefits are primarily associated with releases rather than states, the search process should concentrate primarily on determining appropriate values for the target releases. The state parameters should not be included as parameters of the operating policy for the purposes of search. Rather, these state parameters should be determined in some manner from the transfer function and the target releases.

4) The non-stochastic nature of the linear programming model causes consistent underestimation of near-optimal target values associated with a standard operating policy for a stochastic simulation model.

5) Data structure and data handling capabilities can delimit the feasibility of the various possible search techniques. A difficult-to-handle model does not encourage extensive use of search processes.

6) An adequate linear programming model performs well to select initial values for the search parameters for the more detailed and sophisticated simulation model presented here.

CHAPTER V

THE CONNECTICUT RIVER BASIN MODEL

5.1 Introduction

Mr. David Hellstrom, of the U.S. Army Corps of Engineers, has developed a detailed generalized simulation model for river basins. The model, presently in final stages of development, contains a number of interesting and unique features, and differs significantly in form from the models discussed in Chapters 3 and 4. Consequently, the techniques presented here for optimization of reservoir operation were studied in conjunction with the Hellstrom model to examine in a qualitative manner the particular form that a simulation model should take for application of search techniques for operating rules.

Particular features of the model which make it of interest are:

1) Generality - the model is not developed for a particular basin, and at present can simulate a system containing up to 60 physical control or economic use sites (e.g. reservoirs, power plants, irrigation diversions, etc.)

2) Form of operating rule - The operating rule is based on the assignment of priorities to various uses and storage levels throughout the system, and operates iteratively to balance flows and releases in accordance with the priorities. Thus, the relative desirability of achieving certain goals can be indicated by the priorities assigned.

3) Possibility of Wide implementation - It is expected that the model may be widely used throughout the Corps of Engineers for basin planning studies.

The results presented are limited, due to major revisions continually being made in the form of the model, partly as a consequence of the results of these studies, and computer turn-around time considerations, limiting to one run per day. However, a number of qualitative conclusions will be presented based on the experience with this form of model.

The model is currently being implemented with data from the Connecticut River Basin in New England. For this basin, no mathematical programming model was available to provide a starting point for the search on operating parameters, and the scope of this investigation precluded the development of such a model. As a consequence, the starting point for optimization proceeded from those parameters defining and best simulating the historical operation of the basin over the five years 1957-61.

5.2 The Simulation Model

5.2.1 The Basin

That portion of the Connecticut River Basin which is simulated is shown in Figure 5.1. The 5 upstream reservoirs are used for over-month storage, and the 5 downstream reservoirs are used for pondage. The simulation model operates by months, hence the 5 downstream pondage reservoirs do not appear as reservoirs in the model. Power plants located at Moore and Comerford Dams, and at the 5 downstream sites, where operation is run of the river. Capacities of the facilities are presented in Table 5.1.

Due to computer time limitations, each submittal to the computer was limited to 5 minutes. This allowed for four different parameter sets, each with a simulation run of 5 years, for each submittal. Although synthetic hydrology can be used with the model, it is currently being operated with the historical record for the 5 years 1957-61.

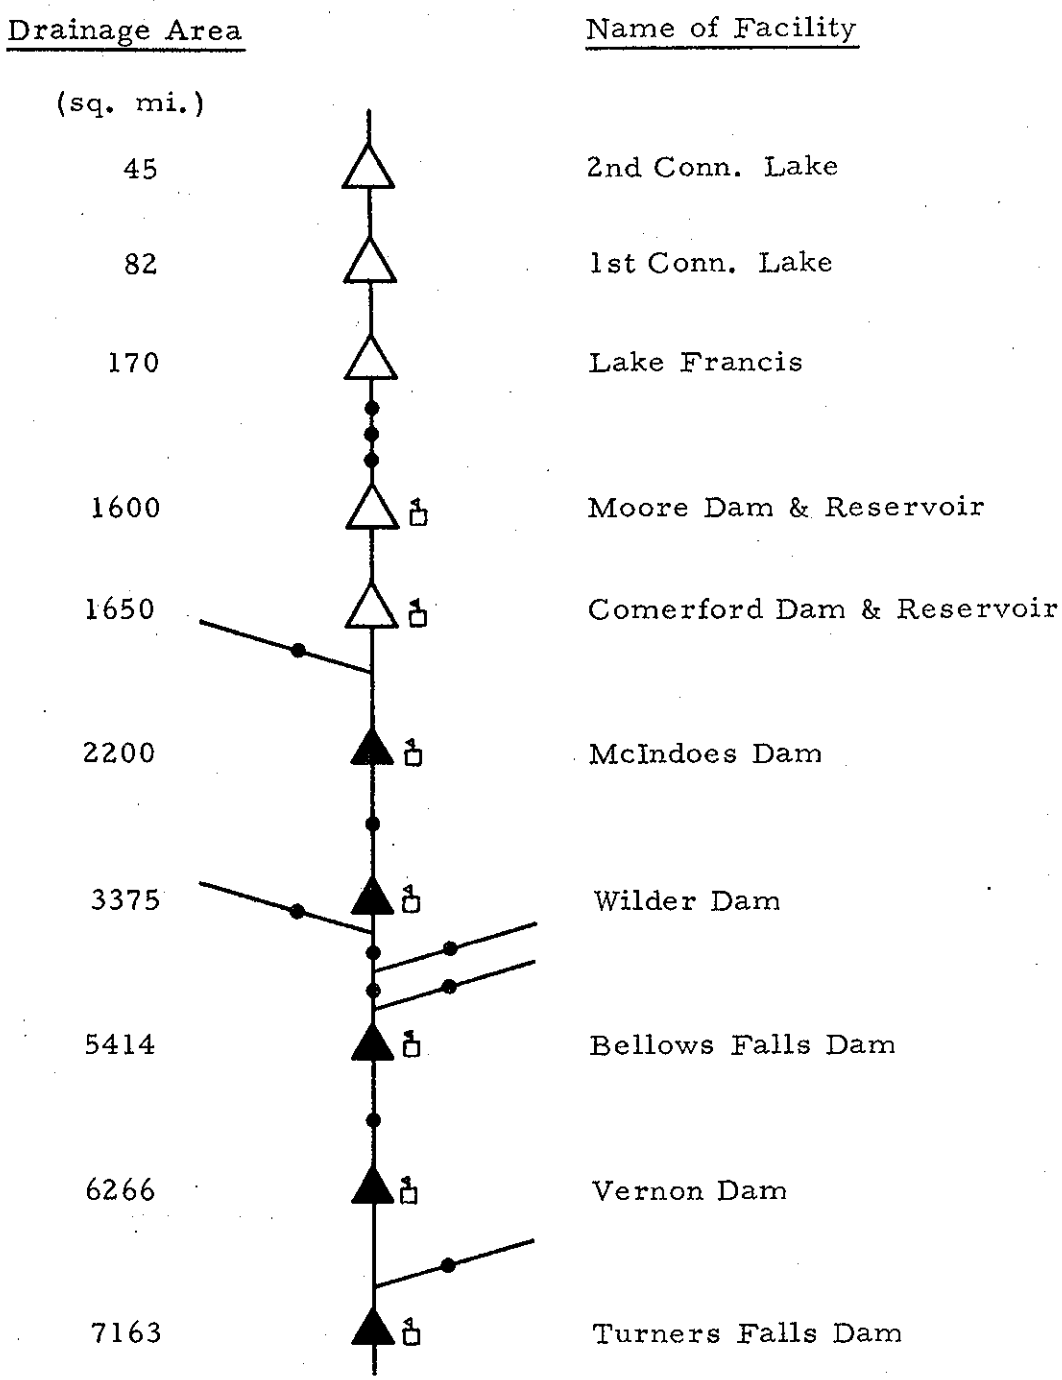

FIGURE 5.1

SCHEMATIC DIAGRAM

CONNECTICUT RIVER BASIN

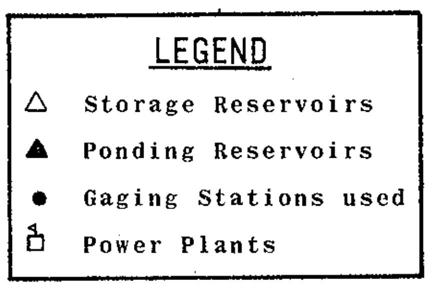

TABLE 5.1

CAPACITIES OF FACILITIES IN THE CONNECTICUT RIVER BASIN

Reservoir	Capacity (month-second-feet)
2nd Connecticut Lake	195
1st Connecticut Lake	1268
Lake Francis	1645
Moore	1890
Comerford	535

Power Plant	Monthly Maximum Energy (MWH/month)
Moore	138800
Comerford	109000
McIndoes	7300
Wilder	24100
Bellows Falls	32850
Vernon	20450
Turners Falls	37200

The basin is presently operated principally by the private power companies for power production. The Corps of Engineers has several flood control reservoirs, but they contain negligible conservation pools and were not included in the simulation for this reason. Power production is primarily peaking power over a base load provided by thermal plants. At present, reasonable benefit data is assumed only for power production, and not for other uses.

In order to test the optimization features, maximization of system power benefits subject to existing power demands was selected as the objective function. Although this selection implies single-purpose operation of the basin, and hence is a special case, some operation for flood control and low-flow augmentation is included in the form of the rule curve constraints for the reservoirs. The single-purpose economic objective was selected because only the values for power benefits were felt to represent reality, and data for other uses was not available. The single-purpose operation does not reduce the problem to a trivial one, however, since the demand varies by months and the benefit curve is non-linear. Thus, the problem becomes one of determining the operation of 5 reservoirs in series to best meet target demands.

5.2.2 The Model

5.2.21 General Structure of the Model

The simulation model is fully documented in reference (18). A brief outline of the main features of the model will be presented here.

The model performs a simulation run with a monthly time increment, generating pertinent physical output at points of interest in the basin. After an entire run has been completed, the physical output is converted to economic benefit through appropriate benefit functions. Thus, the operation of the basin is independent of economic factors, and depends on physical parameters only. Output may consist of a number of specifications of physical and economic data, and the data structure allows for simple modifications. The present form of the program provides for up to 4 different sets of input data for each submittal to the computer.

The program routes mean monthly flows throughout the system, and no flow lag is considered except for mass conservation requirements at storage reservoirs. Evaporation losses are not at present included in the model.

A detailed power subroutine calculates energy generated at a point and in addition the potential energy that could be generated at all sites, whether or not installed capacity exists. Provision is made for considerations of efficiency, dump energy production, tail water elevation, and variation in head at a plant.

5.2.22 The Operating Policy

Although the operating policy is contained in a separate subroutine, the entire structure of the simulation model is closely linked with the structure of the operating rule. Thus, changes in operating policy can come easily only through changes in the parameters of the program, not in the form of the operating rule.

The operation is based on determining the change in storage at each reservoir in the basin for each month. A value for the storage change is arrived at by iteratively balancing out "priorities" which reflect flow conditions throughout the basin. Priorities are obtained as functions of either flow, storage, or energy production at a point of interest, and the shape of the functions can be defined by parameters of the operating rule.

For reservoirs, the priority function is obtained from specification of 6 parameters at each reservoir, 2 of which can vary from month to month. The parameters are defined as follows, using the Fortran mnemonics, where N represents an index for the particular reservoir, and M a monthly index, from 1 to 12:

CAPA(N)	Capacity at spillway crest, hence maximum pool level
CAPB(N,M)	Storage at upper limiting rule curve for month M
CAPC(N,M)	Storage at lower limiting rule curve for month M
CAPD(N)	Storage at top of dead storage pool
PRIUPP(N)	A dimensionless number defining the shape of the priority function in the region between CAPA and CAPB, between 0 and +9.99
PRILOW(N)	A dimensionless number defining the shape of the priority function in the region between CAPC and CAPD, between -9.98 and 0.

These parameters are used to define the reservoir priority function, as shown in Figure 5.2. The region between CAPA and CAPB is called "Flood Control Storage", that between CAPC and CAPD the "conservation storage", and the region between the upper and lower rules curves, CAPB and CAPC, is termed the "neutral zone", and operation of the reservoir in this range is associated with a priority of 0. The reservoir is indifferent to changes in storage so long as it remains in this range. As the storage increases out of the neutral zone, increasing positive priorities are incurred up to a maximum of +9.99, at the spillway crest capacity. No priority can exceed +9.99. Similarly, operation at CAPD is associated with the largest negative priority possible, -9.98.

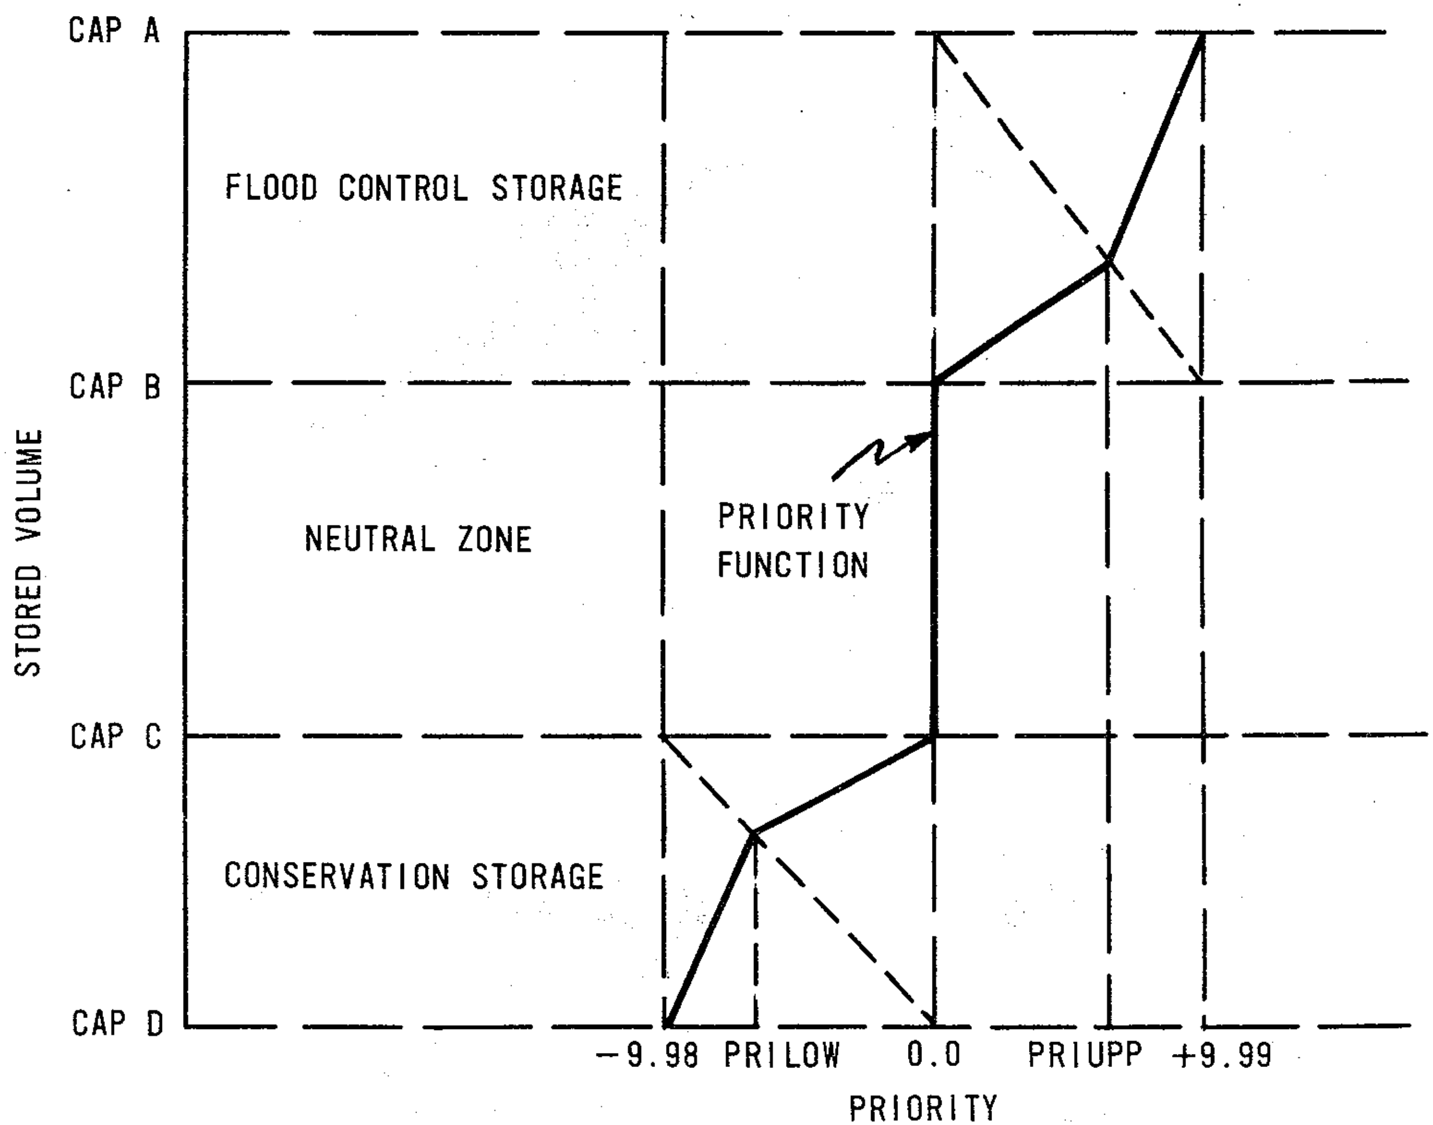

FIGURE 5.2

RESERVOIR PRIORITY FUNCTION

Since the iteration procedure attempts to reduce extreme priorities, the operating procedure attempts to maintain all reservoirs in the neutral zone. The values of PRIUPP and PRILOW serve to determine the severity in terms of incurring priorities of operation outside the neutral zone. Thus, a high value of PRIUPP produces a priority function that increases quite rapidly relative to capacity, that is, small increments in capacity above CAPB produce large increments in priority for the reservoir, and thus are not favored in the iteration process. Similarly, large negative values of PRILOW indicate a reluctance to move into the conservation storage region. Values of the priorities near zero indicate a relative indifference to excursions out of the neutral zone.

A similar function is used for discharge for river points, with zones of flood damage, neutrality, and pollution control, and will not be discussed in detail. Operation is exactly analagous, with 6 parameters being specified at a point.

Energy production priorities may also be included in the iteration process to operate for power production. The installed capacity places a constraint on the maximum energy produced during the month, with operation at 100% load factor. The peaking plants of the Connecticut River Basin operate at much lower load factors. Thus, for each plant, a target energy for each month is established as some percentage of the maximum energy. This target energy will vary from month to month in response to changing demand for peaking power in the regional grid, and is considered to be imposed externally by power system needs, and hence is not a parameter of the operation. That is, although increasing the target energy for the system may yield increased simulated benefits, this would not realistically reflect the example problem at hand, which is to best operate the system to meet existing targets as well as possible.

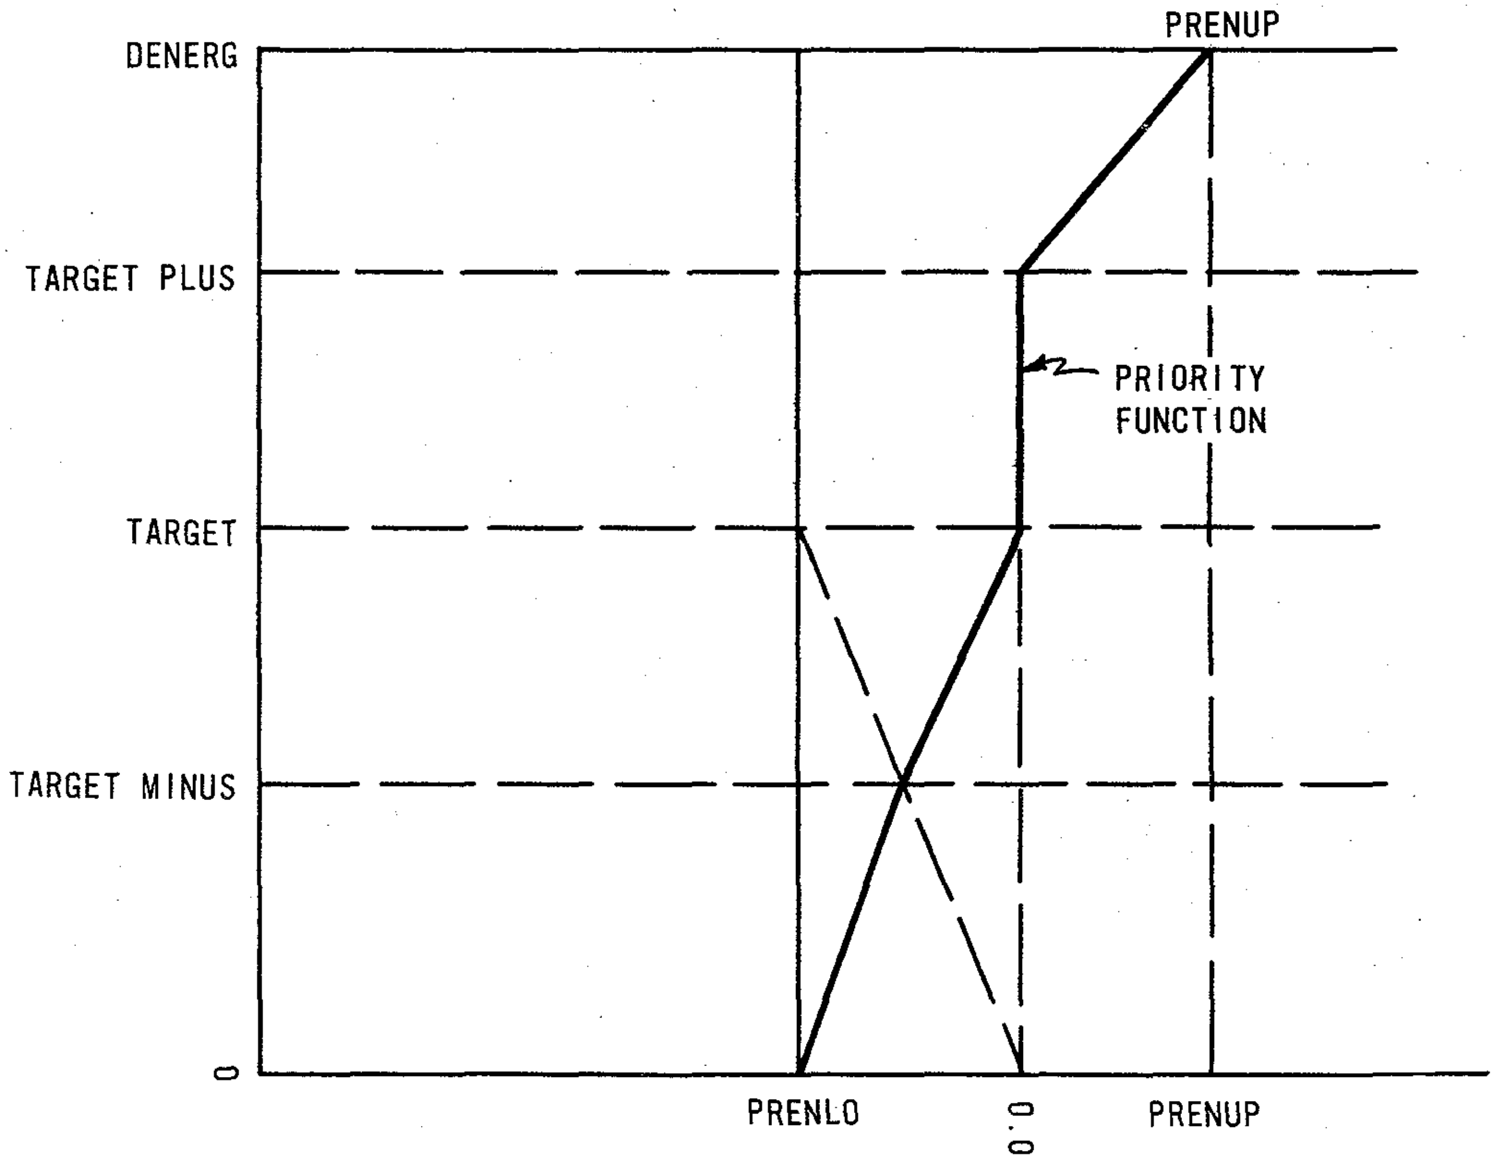

FIGURE 5.3

RESERVOIR PRIORITY FUNCTION

Generation of excess dump energy puts the function in the zone between TARGPL and DENERG, with steadily increasing positive priority, to a maximum of PRENUP at DENERG. Falling short of the target results in an increasing negative priority, to a maximum absolute value of PRENLO at zero energy production. By appropriate selection of the priority values, the priority function can be structured to strongly or weakly favor production of target or dump energy.

With the priority functions as defined above, the iteration procedure used to generate the values of storage change for each reservoir in each month may be briefly described.

For the initial trial in any month, the operating rule maintaing all reservoirs at their levels of the previous month, i.e. no storage change, and routes the incoming river flows through the system. As a results, priorities are incurred at various sites throughout the basin, due to failure to meet targets, excess flows, operation outside rule curve limits, etc.

On succeeding trials, each reservoir attempts to change its storage so as to decrease the absolute value of the largest priority that it can affect, either downstream or that produced by its own storage level. Each reservoir performs this operation by itself, without regard to releases from other reservoirs. For each reservoir, an increment of storage change is determined at the beginning of each month, and the operating policy than either applies this fixed increment to the reservoir storage, either increasing or decreasing the storage by the fixed increment, or maintains storage at the level of the previous iteration.

After all storage changes have been so determined in any given iteration step, a new set of priorities representing the new system conditions is developed. The process then repeats until either no further improvement in terms of reduction of priority is possible, or the maximum allowable number of iterations, a user-defined parameter, is reached. For the situation studied in the model, a maximum of 30 iteration steps was allowed, and the system usually was "balanced" prior to reaching 30 iterations. At the end of each iteration, the increment of storage change for each reservoir is reduced by some fixed percentage, to provide a damping effect on the balancing process by restricting the range of storage for the reservoir. At this point, information on the physical conditions for the month is written on disk storage, and the simulation proceeds to examine the succeeding month.

5.2.23 The Benefit Function

At the end of the entire simulation run, economic benefits are determined as a function of monthly physical outputs, which have been stored by the program. For the single-purpose power benefit study used here, the benefit function is taken as a 2-piece piece-wise linear function as shown in Figure 5.4. Slope BP1 is the rate used to calculate benefits at the inflection point, representing precisely meeting the target demand. Dump energy increases benefits at the lesser rate BP2, up to the maximum energy possible at the site. Falling short of the target requires that the deficit be made up at a higher rate from other sources, hence the net rate BP3 is less than BP1. Thus, the particular shape of the energy benefit function is derived from the three fixed parameters BP1, BP2, and BP3, and the variable target energy for each site for each month, which defines the position of the breakpoint. For the present case, the values of BP1, BP2, and BP3 are taken as 4, 2, and 5 mills per kilowatt-hour respectively.

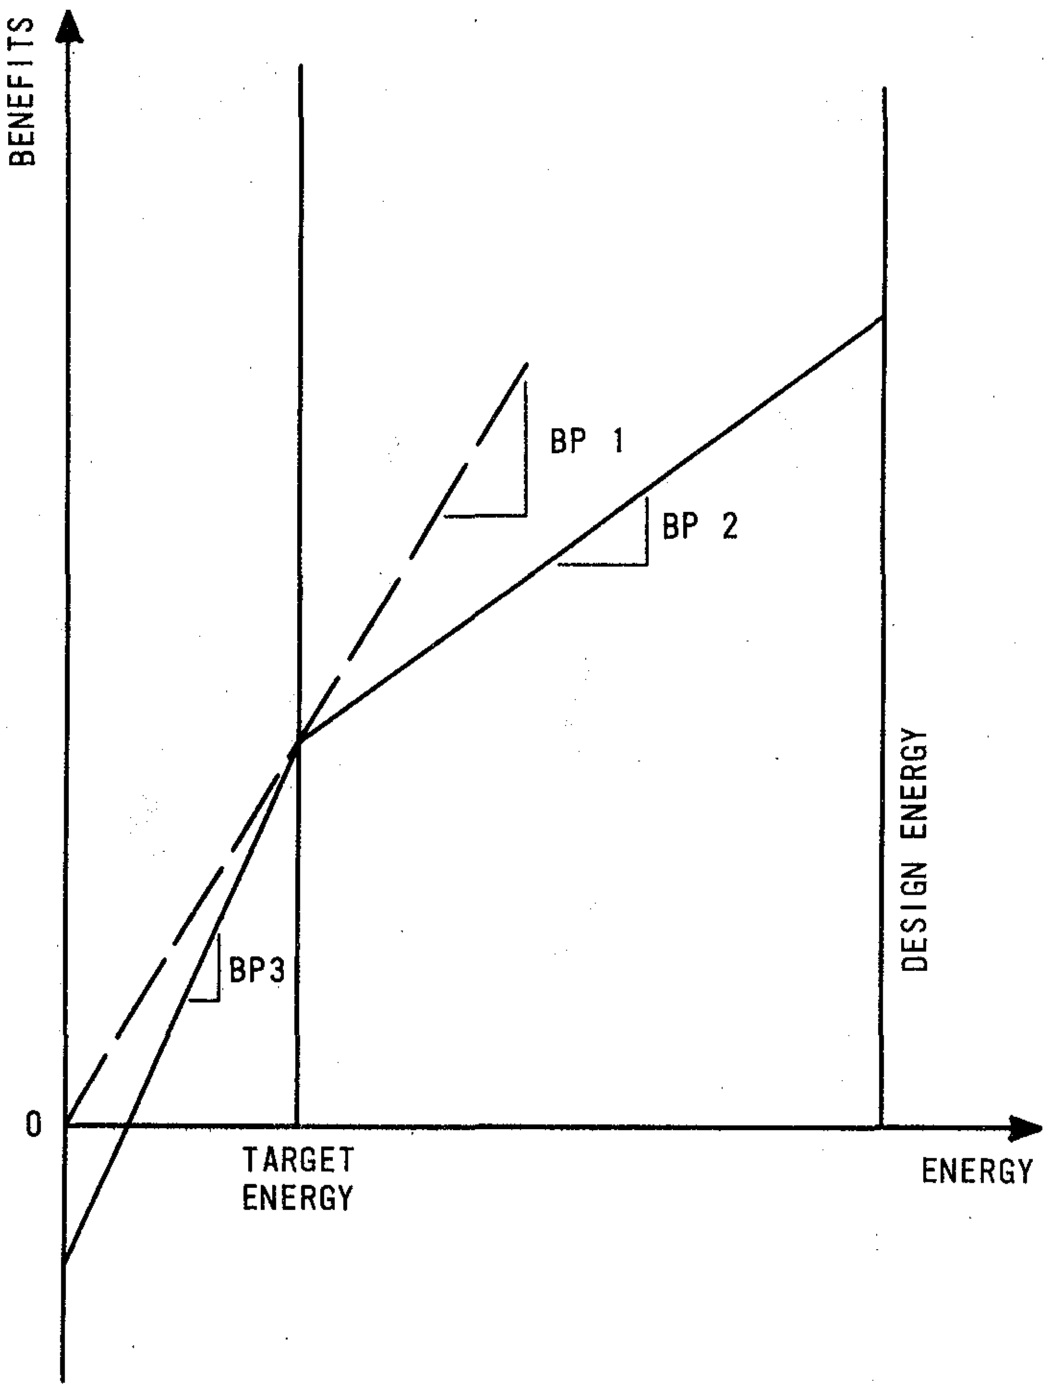

FIGURE 5.4

ENERGY BENEFIT FUNCTION

5.3 Studies with the Model

5.3.1 Phase I

As developed by Hellstrom, the model was used primarily to simulate as well as possible the historical operation of the basin during the years 1957-61. The form of the operating rule and the parameters associated with it were chosen primarily with this purpose in mind, and the simulation performed under these conditions.

For the purposes of the present study, operation that maximized gross power benefits was the objective. The original form of the operating rule, developed to simulate historical flows, did not perform adequately when parameters to maximize power benefits were used. The iteration procedure did not balance the reservoir priorities properly. Consequently, during Phase I the form of the operating rule was revised continually, until the form of operation described in the previous section was developed. Once the form of the operating rule was fixed, search studies were commenced again, and comprise Phase II. Thus, no qualitative results are presented for Phase I, but conclusions are drawn from the experience derived during Phase I studies. Qualitative results are presented in Section 5.4.1

5.3.2 Phase II

Phase II studies consist of a series of experiments to determine appropriate parameters to sample in policy space, and to select an appropriate starting point. In the absence of a mathematical programming model, a pure LPSS technique could not be used.

The search was carried out heuristically after detailed examination of the physical output and performance of the operating rule for each Q set. Experience gained in Phase I studies in identifying pertinent variables was used. Gross economic power benefits were obtained for annual interest rate of 0 and 4%. For the case of 4% interest rate, benefits were discounted to present value at the start of the 5-year simulation. Total power generated was also examined.

The starting point selected initially for Phase II studies was the parameter set which best simulated the historical operation of the basin. This set was tested both with and without power priorities. Heuristic search was carried out from this point to define a near-optimal region, with results presented in Table 5.2a. After the optimal region was defined, studies were carried out to examine in more detail the nature of the near-optimal peak.

5.4 Results

5.4.1 Phase I

Initial studies were made by setting reservoir priorities at low levels and power priorities at high levels, so that reservoirs would respond to power demand only. Results of simulation runs yielded extremely low benefits, and in particular, loss of head at Moore and Comerford Dams was severe.

This pointed out that maximization of power benefits could not be achieved internally by the operating procedure by setting high priorities on power. Investigations showed that the achievement of high head levels at Moore and Cwmerford was one of the major factors in yielding high power benefits, and that an appropriate balance had to be found at these two sites between energy and reservoir priorities. This balance was found to exist away from the regions of extreme priority values. Further results indicated that selection of high power priorities for the smaller downstream run of the river plants often would cause depletion of upstream reservoirs without yielding commensurately greater benefits. Operation of the three upper lakes was also found to be a significant factor in achieving power targets at Moore and Comerford. Thus, Phase I studies pointed out the following results for use in Phase II:

1) Maximum benefits could not be achieved through high power priorities and low reservoir priorities.
2) Achievement of high head levels at Moore and Comerford was a primary physical objective for generation of high power benefit.
3) Power priorities could not be set uniformly throughout the basin, since Moore and Comerford were the most significant producers of energy, and the downstream run of the river plants would draw too much from the upstream reservoirs if power priorities were set on a level with Moore and Comerford.
4) At Moore and Comerford sites, the balance between reservoir and power priorities was significant.
5) The three upper lakes could best be operated to regulate flows for Moore and Comerford reservoirs, responding to downstream power priorities but exerting some damping influence due to moderate priority levels at their own sites.

5.4.2 Phase II

Results for Phase II studies are presented in Tables 5.2a and 5.2b. Power production and economic benefits at 0 and 4% are presented for each simulation run. Pertinent parameters for the run are given in the Q set. Phase I studies served to define the important operating parameters to be included in the Q set. These parameters include the power and energy priorities at Moore and Comerford sites, the shape of the rule curve at these two sites, and the priority values associated with

TABLE 5.2

CONNECTICUT RIVER BASIN RESULTS

Run	Power (mwh x 10^7)	\$ Gross Benefits 0%	\$ Gross Benefits 4%	Moore Power P.	Moore Res. P.	Comerford Power P.	Comerford Res. P.	Lake Francis	First Conn. Lake	Second Conn. Lake	Rule Curve
1	.670	20,474,117	18,211,479	0 0	-6 +1	0 0	-6 +1	-5.3 +1	-5.4 +1	-5.5 +1	historic
2	.667	22,182,996	19,761,393	-8 +8		-8 +8					
3	.670	22,213,677	19,789,342								high head
4	.673	22,267,384	19,384,870	-7 +7		-7 +7					
5	.675	22,233,833	19,801,113	-6 +6		-6 +6					
6	.668	21,999,174	19,612,554	-8 +8	+5	-8 +8	+5	-4			
7	.665	21,786,142	19,421,535					-5			large neutral zone
8	.660	21,397,698	19,064,094								drawdown
9	.673	20,468,710	18,201,055	0 0		0 0					
10	.674	22,308,821	19,872,205	-7 +7		-7 +7		-4			high head
11	.669	22,042,134	19,633,607					-5			large neutral zone
12	.666	21,881,194	19,501,286								drawdown
13	.668	20,268,436	-------	0 0		0 0					

Table 5.2a

Identification of Near-optimal Region

TABLE 5.2 (CONT.)

CONNECTICUT RIVER BASIN RESULTS

Run	Power (mwh x 10^7)	$ Gross Benefits 0%	$ Gross Benefits 4%	Moore Power P.	Moore Res. P.	Comerford Power P.	Comerford Res. P.	Lake Francis	First Conn. Lake	Second Conn. Lake	Rule Curve
14	.675	22,235,744	19,809,014	-7 +7	-6 +1	-7 +7	-6 +1	-2 +1	-2 +1	-2 +1	high hear
15	.674	22,303,447	19,869,680					-3	-3	-3	
16	.674	22,301,652	19,866,720					-4	-4	-4	
17	.673	22,302,669	19,865,984					-5	-5	-5	
18	.673	21,949,127	19,563,131	-6 +6		-6 +6		-3	-3	-3	
19	.672	21,932,823	19,551,154	-6.5 +6.5		-6.5 +6.5					
20	.672	21,917,605	19,538,538	-7 +7		-7 +7		-2.5	-2.5	-2.5	
21	.673	22,027,130	19,633,921					-3.5	-3.5	-3.5	
22	.670	21,958,647	19,579,240	-7.5 +7.5		-7.5 +7.5		-3	-3	-3	
23	.669	21,920,110	19,543,975	-8 +8		-8 +8					
24	.671	21,943,398	19,563,447	-7 +7		-7 +7					
25	.671	21,983,395	19,597,089					-4	-4	-4	

Table 5.2b

Examination of Near-optimal Region

Lake Francis and the first and second Connecticut Lakes. For purposes of simplicity in the tabulation, these values are entered in the tables only when they show a change from a previous value.

The entries under the "Rule Curve" column require some explanation. The historical rule curve is the one actually followed by the power companies in the area, and used in the historical simulation. The rule curve insuring high head levels at Moore and Comerford is obtained by setting CAPA, CAPB, and CAPC within 1 msf of each other, thus insuring that the neutral zone for these reservoirs is at the top of the pool. Lowering the pool by withdrawal incurs a negative priority, tending to force the pool level back up again. To test the sensitivity to rule curve operation, the distance between CAPA, CAPB, and CAPC was set to 50 and 100 msf respectively for Comerford and Moore stations for runs 7 and 11. This serves to create a larger neutral zone, allowing the reservoir more freedom to respond to changing conditions. Another attempt at changing the rule curve is presented in runs 8 and 12. For these runs, the high head rule curve was maintained except for the months of January, February, and March, when the reservoirs normally draw down in actual operation to accept the spring inflows.

Results presented in Table 5.2a delineate the Q set which yields the near-optimal region. Runs 1 and 2 simulate historical operation of the basin, with and without power priorities respectively. Runs 3, 4, and 5 test the high head rule curve with varying priorities at Moore and Comerford power plants. Downstream power priorities are held at values of +3.0 and -3.0 for all run of the river plants, as indicated by Phase I studies, for all the studies of Phase II.

Examination of detailed physical output for the first 5 runs showed that Lake Francis was storing more water than necessary. Therefore, run 6 reduces the value of PRILOW for Lake Francis to -4, making it more responsive to demands for water. Runs 7 and 8 vary the shape of the rule curve at Moore and Comerford stations by providing a larger neutral zone and drawing down for the spring inflows. Run 9 investigates the effects of the drawdown rule curve i power priorities are ignored. Runs 10 through 13 repeat the previous 4 runs with smaller values for energy priority at Moore and Comerford.

Results indicate that the Q set associated with Run 10 yields maximum gross benefits, even though total power production is not as great as that obtained in run 5. Having identified the near-optimal region, studies were undertaken in an attempt to further define this region. Results are presented in Table 5.2b. Runs 14 through 17 test the response to changes in priority values for the upper lakes. Runs 18 through 25 attempt to define appropriate combinations of priorities at Moore, Comerford, and the upper lakes. Examination of the table shows that no improvment on the benefits from run 10 was obtained by the sensitivity studies performed. It is only fortuitous that maximum benefits were obtained in the first stage rather than the second.

As can be noted from an examination of runs 1 and 2, the inclusion of power priorities results in a large increase in power benefits, while causing a drop in total power produced. Further, improvements in gross benefits beyond those developed in run 2 are not large, being limited to an increment of less than 1% of total benefits obtained in run 2.

5.5 Conclusions

In spite of the fact that a formal LPSS process could not be used, simple search techniques over a limited set of operating parameters did result in increases in gross benefit. High benefits due to run 2, the historical operation with power priorities, can be attributed to the fact that the power companies have developed their operating rules by experience over a large number of years, and with extensive hydrologic studies. Moreover, the basin facilities have been designed, and the target power demands selected, to "optimize" use of the available water. Thus, the absence of a mathematical programming model is not a severe drawback in the present case. The construction of a new facility or changing external power demands and hence assoicated target values, or the utilization of a variety of synthetic hydrologies would decrease the value of the historical operation as a starting point, and a mathematical programming model would be appropriate.

An examination of Table 5.2 shows that the benefit surface is fairly flat in the near-optimal region. This observation served to limit the search process, since increases in benefit would be small from additional searching.

For the 5-year period of simulation used in this study, no significant effect of interest rate on selection of the near-optimal Q set was noted. That policy which is near-optimal at 0% is also near-optimal at 4%.

Thus, quantitative results of the study indicate that, with a model of this type, search techniques can yield an improved value of gross benefits. However, qualitative experience with the model did not find it particularly well suited to search optimization.

The parameters of the model are primarily priority levels and state parameters (rule curves). Search studies were carried out over these parameters to determine the near-optimal region. It must be noted, however, that these parameters are removed a number of heirarchical levels from the generation of benefits. Benefits are associated primarily with releases from reservoirs. Thus, an operating policy with parameters which are target releases is one heirarchical level removed from benefits. A policy where the parameters are state parameters first requires evaluation of releases as a function of state parameters, and is consequently two levels removed from benefits. The present policy requires priority parameters, which are used to determine states, and as such is three levels away from production of benefits.

Such a policy proves applicable for simulation studies, since the alrge number of parameters required to go from priorities to physical releases introduces a good deal of flexibility. By determining appropriate values of these paraemters, the model can simulate historical operation quite well. In the Hellstrom model, when it was recognized that maintenance of high head levels at Moore and Comerford was a physical objective, this was easily accomplished since reservoir state parameters were parameters of the operating rule. A release-oriented rule would have to accomplish this on the basis of constraints on reservoir volume, since desired state levels cannot be obtained by setting release parameters.

For optimization studies, however, the fact that operations parameters are three levels removed from benefits makes it difficult to determine appropriate changes in parameters from an examination of benefits. The intermediate stages tend to mask the consequences of any such changes in terms of benefits. The above concepts point out that a non-optimizing operating rule should be parameterized by selecting parameters which are closely associated with benefits or objective goals.

A further consequence of the introduction of heirarchical levels for operating parameters beyond those directly associated with benefits is that the formulation of mathematical programming models is not straightforward. For the present study, it is not at all clear how to formulate a linear programming model to yield values for the priority parameters.

The experience with this portion of the study may be summarized as follows:

1) A modified search technique does lead to improvements in benefits even though the exact nature of the operating parameters is not clear.

2) The efficiency of LPSS techniques is conditioned strongly by the form of the model. Parameters of the operating rule should relate closely to methods of obtaining benefits. Thus, if benefits accrue principally to releases, the parameters should be releases; if benefits accrue principally to the maintenance of certain state levels, then state values are appropriate parameters for the operating rule. As noted both in the present study and in the Maule study, parameters of the operating rule which are non-physical (priorities, shape parameters) cannot be readily associated with mathematical programming models, and appropriate search techniques on such parameters are not obvious. The concept of hierarchy of such parameters in terms of relation to ultimate production of benefits is useful in evaluating simulation models for applicability of LPSS techniques.

CHAPTER VI

CONCLUSIONS AND RECOMMENDATIONS

6.1 Conclusions

The principal conclusions to be drawn from the total investigation are the following:

1) From conceptual analysis of the search for optimal operating rules, it is clear that the relative value of an operating rule depends upon the Q set and that sophisticated optimization techniques may not produce true optimal rules because of the Q set considerations.

2) The proposed linear programming-simulation-search technique for obtaining near-optimal operation of multiple reservoir systems is a feasible method of obtaining "best" estimates of benefits associated with various plans. Use of simple mathematical models and search techniques is sufficient to show significant increases in benefits.

3) The contribution of the linear programming model to the total search technique is much more significant than the contributions of the other search processes.

4) In designing a simulation in which search is to be carried out for operation, the data structure and data handling capabilities are of particular importance in that they determine the amount of search that is feasible.

5) Search on parameters heirarchically removed from the production of benefits is not an efficient process.

More specific conclusions dealing with particular areas of investigation are:

1) Linear programming models consistently underestimate target releases for stochastic models when constrained by availability of water.

2) State-oriented models display great sensitivity to values for the steady-state available water parameters, which must be consistent with the target releases.

3) State parameters should not be included in the search process unless benefits accrue directly to achievement of certain states.

4) The number of periods used in a linear programming model should not be less than the number of periods per year in the simulation model.

5) Historical experience in a given basin may be near-optimal operation.

6) Results of search using the three simulation models did not show sensitivity to interest rate, and this appears to be a limitation in the work presented here.

6.2 Recommendations

Experience with the LPSS process has pointed out certain areas where extensions are desirable. In particular, methods are sought to develop consistent steady-state water parameters for any set of target releases. Linear programming models develop consistent values for the two different types of parameters, but the concept of "steady-state" available water is not directly applicable to the stochastic simulation case. Two methods are suggested below.

The first method of deriving consistent state parameters from the release parameters is based on statistical studies using the simulation model. A release-oriented policy using the desired target parameter set can be used in the simulation model. Mean values of water available can then be calculated statistically from the detailed physical behavior of the system. These values can then be used in association with the target parameters in a state-oriented policy. Mean available water can then be determined statistically for results of a simulation with this operating policy, and the state parameters revised again. If desired, the process can be repeated until the state parameters "reproduce" themselves statistically.

The above method is based on the statistical steady-state concept for the simulation model. A statistical steady-state is the only one possible for stochastic models, An alternative approach is to use the steady-state equations which define the linear programming model. These equations can be solved with the values of the target releases set at desired levels. Consistent steady-state available water values are then obtained. For this case, the linear programming model is used as the transfer function defining system response.

The consistent underestimation of target releases by linear programming models suggested that the linear programming model be constructed with additional parameters which would allow modification of the nature of the results. Thus, with a programming model for operation only, increasing the inflow by a multiplicative factor may offset the tendency of the model to underestimate, and the use of the linear programming model assures internal consistency of the parameter values.

The above suggests an alternate search method which would insure internal consistency of the operating policy parameters. Rather than search directed to values of these parameters, the search may be performed over the parameters of the Q set. The linear programming model is looked on as a sub-optimization procedure which yields, for any given Q set, the near-optimal values of the operating parameters. If the linear programming model can be characterized by fewer variables than the simulation operating policy, then the search may be more efficient. For example, a multi-reservoir monthly simulation may involve a large number of target releases. If, in a linear programming model, parameters such as total power export target or total irrigated farming target can be included, then the search may be performed over values of these targets as independent variables, with the operating policy targets being determined as dependent variables by the linear programming model.

As noted in Chapter 3, comparisons of the simulation and linear programming model can yield valuable information. A disparity in the Q set or other parameters can prevent such a comparison. Option should be provided in the simulation for making the Q set identical to that in the linear programming model.

6.3 Suggestions for Further Work

Appropriate extensions of the methods described here have been indicated in the previous section and throughout the text. Experiences with the three simulation models pointed out that data handling and data structure was a major difficulty in using the models, particularly for search techniques. The generation of appropriate data structures and data handling methods for the various models is seen as a vital contribution to the state of the art. The development of such data structures has been a prime concern for structural and transportation systems studies, and has been seriously neglected for water resource systems.

The generation of an appropriate data structure for the models will require study of the nature of the simulation and mathematical programming models, and the type of consideration to be made of such features as power, irrigation, and flood control use. A number of possibilities are available, and an investigation into the merits of these possibilities would be a great service to the systems planner.

The data handling problem is essentially a systems programming problem, but it is clearly related to the data structure problem. Ease of specification, ease of modification, and selective output would be important considerations. Once such a data handling procedure is developed, the possibility of direct interfacing between the simulation, the linear programming model, and the search process becomes available. Such interfacing would allow for detailed search techniques of the type indicated in Chapter 4, even for relatively complex models.

A problem-oriented language for searching the policy space could profitably be developed, with the aid of a system such as ICES (22), which allows the user to define commands and procedures pertinent to the problem at hand, after the data structure and handling problem is solved. Typical commands would allow for simple specification of a uniform grid search by specification of the ranges of the parameters and the grid spacing. Thus a command of the form:

UNIFORM GRID X 10. 20. Y 8. 13. SPACE 1.

would specify in this simple manner that the search points were to be carried out on a uniform two-dimensional grid for the variables X and Y, X ranging from 10 to 20, and Y from 8 to 13 units, with a grid spacing of 1 unit, for a total of (11 x 6) = 66 points in the grid.

A command of the form:

SEARCH SAVE BEST

would initiate a simulation for each of the points of the uniform grid. The values pertaining to the best point would then be stored in the computer. Similar commands for single-factor and marginal analysis could be utilized to rapidly yield a search procedure suitable to the particular problem. Pre-programmed optimal search techniques could then be profitably developed and utilized.

When the above suggestions have been implemented, a powerful tool for investigations of water resources systems will be in the hands of the engineering decision-maker. In the author's opinion, such developments are the most fruitful line of investigation for those engaged in water resource systems research. Such a tool would then be able to make optimal use of the sophisticated optimizing procedures and search techniques which are more properly the field of systems analysts, mathematicians, and specialists in operations research, but would be oriented primarily towards utility for decision-making in systems planning in light of engineering reality and current theoretical capabilities.

REFERENCES

(1) Maass, A., et al, Design of Water Resource Systems, Harvard University Press, Cambridge, 1962

(2) Hufschmidt, M., Chapter 10, Design of Water Resource Systems, ibid.

(3) Dorfman, R., Chapter 13, Design of Water Resource Systems, ibid.

(4) Thomas, H., and Watermeyer, P., Chapter 14, Design of Water Resource Systems, ibid.

(5) Wallace, J., "Linear Programming Analysis of River Basin Development", Unpublished Ph.D. Thesis, Department of Civil Engineering, Massachusetts Institute of Technology, Cambridge, Mass., June 1966

(6) Parikh, S.C., "Linear Dynamic Decomposition Programming of Optimal Long-Range Operation of a Multiple Multi-purpose Reservoir System", Proc. Fourth International Conf. on Operations Research, Boston, Massachusetts, September 1966

(7) Hall, Butcher, and Esogbue, "Optimization of the Operation of a Multiple-Purpose Reservoir by Dynamic Programming", Water Resources Research, Volume 4, Number 3, June 1968

(8) Bower, B., Hufschmidt, M, and Reedy, W., Chapter 11, Design of Water Resource Systems, op.cit.

(9) Hufschmidt, M., and Fiering, M., Simulation Techniques for Design of Water-Resource Systems, Harvard University Press, Cambridge, 1966.

(10) Schweig, Z., and Cole, J.A., "Optimal Control of Linked Reservoirs", Water Resources Research, Volume 4, Number 3, June 1968

(11) Little, J.D.C., "The Use of Storage Water in a Hydroelectric

System", Journal of the Operations Research Society of America, Vol. 2, No. 3, pp 187-197, 1955

(12) Buras, N., "Conjunctive Operation of Dams and Aquifers", proc. ASCE, Vol. 89, Journal of the Hydraulics Division, HY6, pp 111-131, November 1963.

(13) Males, R., "A Dynamic Programming Model for Reservoir Simulation", Unpublished M.S, Thesis, M.I.T. Dept. of Civil Engineering, Sept. 1964

(14) Young, G.K., "Techniques for Finding Reservoir Operating Rules", Ph.D. thesis, Harvard University, Division of Engineering and Applied Sciences, November, 1966

(15) Ibanez, F, "Operating Policies in the Simulation of River Basin Development" unpublished M.S. thesis, M.I.T. Dept. of Civil Eng., Feb. 1967

(16) Leonvendagar, M., and McLaughlin, R.T., "A Simulation Model of the Upper Maule River Basin", M.I.T. Department of Civil Engineering, in progress

(17) IBM, "1130 Linear Programming-Mathematical Optimization Subroutine System Program Reference Manual", 1967

(18) Hellstrom, D., "Mathematical Program for Simulation of River Basin Systems, in progress, New England Division, Corps of Engineers

(19) Fiering, M., "Queuing Theory and Simulation in Reservoir Design", Journal of the Hydraulics Division, ASCE, Vol. 87, no. HY6, Nov. 1961

(20) Poblete, J.A., "Mathematical Programming Applied to a Water Resource System", M.S. thesis, unpublished, Dept. of Civ. Eng., M.I.T., Sept. 1968

(21) Kleinmann, P., and Torres, J., "Estudio Hidrologico Laguna Maule", Direccion de Riego, MOP, Santiago, Chile, 1961

(22) Roos, D., ICES System Design, M.I.T. Press, Cambridge, 1966

BIOGRAPHY

Richard Michael Males was born in Macon, Georgia, in 1943, and raised in New York City, where he attended the Bronx High School of Science before entering M.I.T. in 1959. In 1963 he received his S.B. degree from the Department of Civil Engineering and entered the graduate school, pursuing a course of studies in water resource systems analysis. He received his M.S. degree in 1964, then entered the Peace Corps, serving in Lima, Peru teaching computer programming and operations research at various universities. In 1966 he was admitted to M.I.T. as a doctoral candidate in the field of water resource systems. He is a member of Sigma Xi, Chi Epsilon, and Tau Beta Pi honorary fraternities.

APPENDIX A - LIST OF FIGURES

APPENDIX B - LIST OF TABLES

Appendix C

Hydrologies for Maule Basin Studies

$(10^6\ m^3)$

Hydrology 1

List of Average Monthly Inflows

Month	Reservoir 1	Reservoir 2	Reservoir 3
1	59.048	22.494	53.866
2	57.690	31.416	154.167
3	61.076	34.541	149.918
4	56.612	28.383	137.673
5	59.015	29.547	121.837
6	64.087	29.999	123.628
7	120.059	49.147	222.363
8	194.595	76.463	310.457
9	171.970	57.856	256.924
10	119.456	31.192	140.939
11	83.664	22.090	84.950
12	69.050	18.314	59.228
Sum	1116.322	431.442	1815.949

Total = 3363.713

Hydrology 2

List of Average Monthly Inflows

Month	Reservoir 1	Reservoir 2	Reservoir 3
1	61.161	23.561	56.872
2	62.234	33.588	152.973
3	63.465	36.791	150.025
4	61.011	31.417	128.090
5	62.531	31.656	124.606
6	65.302	32.276	138.386
7	115.347	47.464	222.584
8	188.199	72.645	319.410
9	192.492	65.751	266.447
10	134.468	35.064	142.650
11	95.576	24.882	87.794
12	75.198	19.301	60.779
Sum	1176.985	454.395	1850.615

Total = 3481.996

Appendix C

(cont.)

Hydrology 3

List of Average Monthly Inflows

Month	Reservoir 1	Reservoir 2	Reservoir 3
1	62.814	22.835	48.161
2	60.055	31.395	124.225
3	65.821	36.655	142.184
4	63.033	36.801	126.839
5	63.502	32.866	120.716
6	64.041	32.039	132.740
7	112.419	47.861	219.141
8	191.765	74.783	278.942
9	169.956	54.940	209.433
10	124.477	32.504	108.483
11	88.090	22.758	68.160
12	70.898	18.441	50.162
Sum	1136.871	443.877	1629.186

Total = 3209.934

Hydrology 4

List of Average Monthly Inflows

Month	Reservoir 1	Reservoir 2	Reservoir 3
1	62.872	23.110	39.707
2	65.230	34.640	103.696
3	63.899	37.026	126.059
4	58.854	32.253	112.247
5	59.142	29.415	114.723
6	63.089	32.063	123.518
7	113.684	47.580	205.813
8	213.012	83.559	297.573
9	182.976	59.415	235.315
10	130.670	34.527	118.310
11	91.256	23.473	76.687
12	72.008	18.311	53.676
Sum	1176.690	455.371	1607.324

Total = 3239.386